

Decomposing Global Feature Effects Based on Feature Interactions

Julia Herbinger^{1,2}

JULIA.HERBINGER@GMAIL.COM

Marvin N. Wright^{3,4,5}

WRIGHT@LEIBNIZ-BIPS.DE

Thomas Nagler^{1,2}

NAGLER@STAT.UNI-MUENCHEN.DE

Bernd Bischl^{1,2}

BERND.BISCHL@STAT.UNI-MUENCHEN.DE

Giuseppe Casalicchio^{1,2}

GIUSEPPE.CASALICCHIO@STAT.UNI-MUENCHEN.DE

¹ Department of Statistics, LMU Munich, Munich, Germany

² Munich Center for Machine Learning (MCML), Munich, Germany

³ Leibniz Institute for Prevention Research and Epidemiology, Bremen, Germany

⁴ Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany

⁵ Department of Public Health, University of Copenhagen, Copenhagen, Denmark

Editor: Pradeep Ravikumar

Abstract

Global feature effect methods, such as partial dependence plots, provide an intelligible visualization of the expected marginal feature effect. However, such global feature effect methods can be misleading, as they do not represent local feature effects of single observations well when feature interactions are present. We formally introduce *generalized additive decomposition of global effects* (GADGET), which is a new framework based on recursive partitioning to find interpretable regions in the feature space such that the interaction-related heterogeneity of local feature effects is minimized. We provide a mathematical foundation of the framework and show that it is applicable to the most popular methods to visualize marginal feature effects, namely partial dependence, accumulated local effects, and Shapley additive explanations (SHAP) dependence. Furthermore, we introduce and validate a new permutation-based interaction detection procedure that is applicable to any feature effect method that fits into our proposed framework. We empirically evaluate the theoretical characteristics of the proposed methods based on various feature effect methods in different experimental settings. Moreover, we apply our introduced methodology to three real-world examples to showcase their usefulness.

Keywords: interpretable machine learning, feature interactions, partial dependence, accumulated local effect, SHAP dependence

1. Introduction

Machine learning (ML) models are increasingly used in various application fields such as medicine (Shipp et al., 2002) or social sciences (Stachl et al., 2020). Their exceptional predictive performance stems from their ability to capture complex non-linear relationships and feature interactions in the data. However, this ability also complicates explaining the inner workings of an ML model. A lack of explainability might hurt trust or might even be a deal-breaker for high-stakes decisions (Lipton, 2018). Hence, ongoing research on model-agnostic interpretation methods to explain any ML model has grown quickly in recent years.

One promising type of explanation is produced by feature effect methods, which explain how features influence the model predictions similar to interpreting the coefficients of a linear model or visualizing the non-linear effects of a generalized additive model (GAM) through splines. We distinguish between local and global feature effect methods. Local feature effect methods—such as individual conditional expectation (ICE) curves (Goldstein et al., 2015) or Shapley values / Shapley additive explanations (SHAP) (Štrumbelj and Kononenko, 2014; Lundberg and Lee, 2017)—explain how each feature influences the prediction of a single observation. In contrast, global feature effect methods explain the general model behavior. Since global feature effects of ML models are often non-linear, they are usually represented as marginal curves instead of single effect numbers; the most popular variants are partial dependence (PD) plots (Friedman, 2001), accumulated local effects (ALE) plots (Apley and Zhu, 2020), or SHAP dependence (SD) plots (Lundberg et al., 2020). These global curves are expressible via an aggregation of local counterparts (e.g., ICE curves are aggregated to a PD curve). However, such an aggregation might cause information loss due to heterogeneity in local feature effects (e.g., see Figure 1). This so-called aggregation bias is usually caused by feature interactions learned by the ML model, leading to a global feature effect that does not appropriately represent many individual observations in the data (Herbinger et al., 2022; Mehrabi et al., 2021). We refer to this heterogeneity as *interaction-related heterogeneity*. Consequently, global explanations can be misleading and cannot give a complete picture when feature interactions are present, e.g., see how the different shapes of ICE curves are not well represented by the global PD curve in Figure 1. This issue is particularly relevant when ML models are trained on biased data, as the model may, in turn, learn these biases (Mehrabi et al., 2021). Although local explanations could reveal the learned bias, it remains hidden in global explanations due to aggregation (e.g., see the COMPAS example in Section 8).




Figure 1: Left: ICE and global PD curves of feature hr (hour of the day) of the bikesharing data set (James et al., 2022). Right: ICE and regional PD curves of hr depending on feature $workingday$. The feature effect of hr on predicted bike rentals is different on working days compared to non-working days, which is due to aggregation not visible in the global feature effect plot (white curve on the left).

To bridge the gap between local and global effect explanations, so-called subgroup or regional explanations have recently been introduced (e.g., Hu et al., 2020; Molnar et al., 2023; Scholbeck et al., 2024; Herbinger et al., 2022). However, several pitfalls and difficulties arise in the realm of regional explanations, e.g., the REPID approach by Herbinger et al.

(2022) is limited to only PD plots and only to one feature of interest (see Section 3 for details). Another common pitfall is the incomprehensibility in high-dimensional settings (Molnar et al., 2022), as it would require analyzing hundreds of feature effect plots, making it difficult for the user to understand the underlying relationships effectively.

Contributions. We introduce GADGET, a novel framework, which provides regional explanations based on feature effects by partitioning the feature space into interpretable subspaces. We prove that GADGET minimizes feature interactions for a predefined set of features using any feature effect method satisfying the *local decomposability* axiom (Section 4.1 and 4.2). We show that popular feature effect methods (PD, ALE, and SD) satisfy the *local decomposability* axiom (Section 4.3-4.5) and illustrate their applicability within GADGET in Section 4.6 with instructions on how to estimate and visualize the corresponding regional feature effects in Appendix C.4. Moreover, we propose several measures to quantify feature interactions based on GADGET, which provide more insights into the learned effects and the remaining heterogeneity related to the interaction (Section 4.7). Furthermore, we introduce and validate the permutation-based interaction detection (PINT) procedure to detect feature interactions in a model-agnostic fashion based on the underlying feature effect method (Section 5). We empirically evaluate the theoretical characteristics of the different methods based on several simulation settings (Sections 6 and 7) and illustrate their usefulness for two real-world examples in Section 8. We provide a third higher-dimensional example in Section 9 and also propose a general filtering procedure for such scenarios. All proposed methods and reproducible scripts for the experiments are available online via <https://github.com/JuliaHerbinger/gadget/>.

2. Background

2.1 General Notation

We consider a feature space $\mathcal{X} \subseteq \mathbb{R}^p$ and a target space \mathcal{Y} (e.g., $\mathcal{Y} = \mathbb{R}$ for regression tasks). The random variables for the features are denoted by $X = (X_1, \dots, X_p)$ and Y for the target variable. The realizations of X and Y are sampled i.i.d. from an unknown joint probability distribution $\mathbb{P}_{X,Y}$ and are denoted by $\mathcal{D} = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^n$. We denote by $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_p^{(i)})^T$ the feature values of the i -th observation. Arbitrary observations \mathbf{x} and random variables X we index as \mathbf{x}_W and X_W (with $W \subseteq \{1, \dots, p\}$) to refer to the feature values and random variables of the subset. Respectively, \mathbf{x}_j and X_j refer to the j -th feature. Negated sets $-W$ (and indices such as in \mathbf{x}_{-W}) denote the set complement. We assume that a true function f maps \mathcal{X} to \mathcal{Y} up to some label noise, e.g., for regression $Y = f(X) + \epsilon$. Under (independent) Gaussian noise ϵ , or equivalently L2-loss minimization f will be the conditional mean $f(\mathbf{x}) = \mathbb{E}[Y|X = \mathbf{x}]$. In supervised ML, we strive to approximate this relationship by a prediction model \hat{f} learned on \mathcal{D} .

2.2 Definition of Interactions and the Functional ANOVA Decomposition

Friedman and Popescu (2008) defined the presence of an interaction between two features \mathbf{x}_j and \mathbf{x}_k by $\mathbb{E}[\partial^2 \hat{f}(X) / (\partial X_j \partial X_k)]^2 > 0$. This condition implies that when the feature \mathbf{x}_j does not interact with any other feature in \mathbf{x}_{-j} , the prediction function $\hat{f}(\mathbf{x})$ can be

decomposed into two distinct functions: $h_j(\mathbf{x}_j)$, which solely depends on \mathbf{x}_j , and $h_{-j}(\mathbf{x}_{-j})$, which solely depends on the remaining features \mathbf{x}_{-j} , i.e.: $\hat{f}(\mathbf{x}) = h_j(\mathbf{x}_j) + h_{-j}(\mathbf{x}_{-j})$.

This definition of feature interactions is closely related to the concept of **functional ANOVA**¹ decomposition, which has already been studied by Sobol’ (1993); Stone (1994) to explain black box functions. The functional ANOVA decomposition aims to decompose a square-integrable prediction function $\hat{f}(\mathbf{x})$ into a sum of lower-order functions:

$$\hat{f}(\mathbf{x}) = g_0 + \sum_{j=1}^p g_j(\mathbf{x}_j) + \sum_{j \neq k} g_{jk}(\mathbf{x}_j, \mathbf{x}_k) + \dots + g_{12\dots p}(\mathbf{x}) = g_0 + \sum_{k=1}^p \sum_{\substack{W \subseteq \{1, \dots, p\}, \\ |W|=k}} g_W(\mathbf{x}_W), \quad (1)$$

where g_0 is a constant, $g_j(\mathbf{x}_j)$ denotes the main effect of the j -th feature, $g_{jk}(\mathbf{x}_j, \mathbf{x}_k)$ is the pure two-way interaction effect between features \mathbf{x}_j and \mathbf{x}_k , and so forth. The last component $g_{12\dots p}(\mathbf{x})$ always allows for an exact decomposition.

The standard functional ANOVA decomposition by Hooker (2004) assumes independent features and ensures that the component functions $g_W(\mathbf{x}_W)$ can be optimally and uniquely defined by imposing the so-called *vanishing condition*² (Li and Rabitz, 2012; Rahman, 2014). This decomposition is widely used in sensitivity analysis to further uniquely decompose the variance of each component from Eq. (1) and derive a feature importance measure (see Sobol indices in Sobol’, 1993; Owen, 2013). Hooker (2007) introduced the generalized functional ANOVA decomposition with a *relaxed vanishing condition*² that leads to a unique decomposition in case of dependent features. These conditions guarantee orthogonality between the components representing different feature subsets and ensure that each component is distinct without redundancy. In this paper, we do not attempt to compute the exact decomposition presented in Eq. (1); instead, we rely on the existence of such a decomposition for our theoretical proofs in Section 4. We refer to Appendix A for more details on the functional ANOVA decomposition and related research regarding its estimation.

2.3 Feature Effect Methods

Below, we summarize global feature effect methods used in our framework in Section 4.2.

Partial Dependence. The *PD plot* introduced by Friedman (2001) visualizes the marginal effect of a set of features $W \subseteq \{1, \dots, p\}$ for a fitted model \hat{f} by integrating over the joint distribution over all other features in $-W$. The number of features in W is typically one or two. The theoretical PD function is defined by³

$$f_W^{PD}(\mathbf{x}_W) = E_{X_{-W}}[\hat{f}(\mathbf{x}_W, X_{-W})] = \int \hat{f}(\mathbf{x}_W, \mathbf{x}_{-W}) d\mathbb{P}_{X_{-W}}(\mathbf{x}_{-W}), \quad (2)$$

-
1. This should not be confused with the **classical ANOVA (Analysis of Variance)**, which is a statistical method that tests for significant differences among group means in a population.
 2. See Appendix A for more details.
 3. The notation here is used consistently with Friedman and Popescu (2008). Please note: While the right-hand side of the definition always refers to the ML model \hat{f} , the f in f_W^{PD} does not refer to the true relationship (see Section 2.1); f_W^{PD} is rather to be read as a new symbol. In the same manner, the “hat” of \hat{f}_W^{PD} is not introduced to refer to \hat{f} now – but expresses the fact that we now estimate based on data.

which is estimated using Monte-Carlo integration: $\hat{f}_W^{PD}(\tilde{\mathbf{x}}_W) = \frac{1}{n} \sum_{i=1}^n \hat{f}(\tilde{\mathbf{x}}_W, \mathbf{x}_{-W}^{(i)})$ at a specific grid point $\tilde{\mathbf{x}}_W$. For constructing a PD curve, we typically use m grid points⁴ denoted by $\tilde{\mathbf{x}}_W^{(1)}, \dots, \tilde{\mathbf{x}}_W^{(m)}$ and visualize the pairs $\{(\tilde{\mathbf{x}}_W^{(k)}, \hat{f}_W^{PD}(\tilde{\mathbf{x}}_W^{(k)}))\}_{k=1}^m$. Goldstein et al. (2015) introduced *ICE curves*, which, for an observation \mathbf{x} , fixes the values \mathbf{x}_{-W} to the observed ones, and varies the values \mathbf{x}_W yielding the ICE function $\hat{f}_W^{\mathbf{x}}(\tilde{\mathbf{x}}_W) = \hat{f}(\tilde{\mathbf{x}}_W, \mathbf{x}_{-W})$, where the superscript \mathbf{x} denotes the observation whose values \mathbf{x}_W are replaced with $\tilde{\mathbf{x}}_W$. ICE plots visualize the pairs $\{(\tilde{\mathbf{x}}_W^{(k)}, \hat{f}_W^{\mathbf{x}(i)}(\tilde{\mathbf{x}}_W^{(k)}))\}_{k=1}^m$, representing prediction changes of each observation as the values of features in W vary. Notably, the PD curve is the point-wise average over all the ICE curves, i.e., $\hat{f}_W^{PD}(\tilde{\mathbf{x}}_W) = \frac{1}{n} \sum_{i=1}^n \hat{f}_W^{\mathbf{x}(i)}(\tilde{\mathbf{x}}_W)$. ICE curves can reveal heterogeneous effects caused by interactions between features \mathbf{x}_W and other features, which remain hidden in the PD curve due to averaging (see Figure 1). However, they do not reveal which features interact with \mathbf{x}_W . In addition, if the model contains complex interactions, it becomes more difficult to identify clear patterns with interacting features (even if we color-code ICE curves as in Figure 1).

ALE. The PD function uses the marginal distribution and is likely to *extrapolate* in empty or sparse regions of the feature space when features are correlated. This means that the prediction function is evaluated at unrealistic combinations of feature values in Eq. (2), which can lead to inaccurate PD estimates, especially for models with high flexibility in capturing complex feature interactions, such as neural networks (Apley and Zhu, 2020). Hence, ALE uses the conditional distribution to avoid extrapolation. The uncentered ALE $f_W^{ALE}(x)$ for $|W| = 1$ at feature value $x_W \sim \mathbb{P}_{X_W}$ and with $z_0 = \min(\mathbf{x}_W)$ is defined by

$$\begin{aligned}
 f_W^{ALE}(x_W) &= \int_{z_0}^{x_W} \mathbb{E} \left[\frac{\partial \hat{f}(X)}{\partial X_W} \middle| X_W = z_W \right] dz_W \\
 &= \int_{z_0}^{x_W} \int \frac{\partial \hat{f}(z_W, \mathbf{x}_{-W})}{\partial z_W} d\mathbb{P}_{X_{-W}|X_W=z_W}(\mathbf{x}_{-W}|z_W) dz_W.
 \end{aligned} \tag{3}$$

ALE first calculates local derivatives weighted by the conditional distribution $\mathbb{P}_{X_{-W}|X_W=z_W}$ and then accumulates the local effects to generate a global effect curve. For interpretability reasons, ALE curves are usually centered by the average of the uncentered ALE curve to obtain $f_W^{ALE,c}(x_W) = f_W^{ALE}(x_W) - \int f_W^{ALE}(\mathbf{x}_W) d\mathbb{P}_{X_W}(\mathbf{x}_W)$.

Eq. (3) is usually estimated by splitting the value range of \mathbf{x}_W in intervals and calculating the partial derivatives for all observations within each interval. The partial derivatives are estimated by the differences between the predictions of the upper ($z_{k,W}$) and lower ($z_{k-1,W}$) bounds of the k -th interval for each observation. The accumulated effect up to the feature value x_W is then calculated by summing up the average partial derivatives (weighted by the number of observations $n(k)$ within each interval) over all intervals until the interval that includes x_W , which is denoted by $k(x_W)$:

$$\hat{f}_W^{ALE}(x_W) = \sum_{k=1}^{k(x_W)} \frac{1}{n(k)} \sum_{i: \mathbf{x}_W^{(i)} \in]z_{k-1,W}, z_{k,W}] } \left[\hat{f}(z_{k,W}, \mathbf{x}_{-W}^{(i)}) - \hat{f}(z_{k-1,W}, \mathbf{x}_{-W}^{(i)}) \right]. \tag{4}$$

4. Instead of using all observed feature values \mathbf{x}_W , an equidistant grid or a grid based on randomly selected feature values or quantiles of \mathbf{x}_W are commonly chosen (Molnar et al., 2022).

SHAP Dependence. Shapley values (Shapley, 1953) have been adapted from game theory to ML to quantify feature effects on a local level by fairly distributing the prediction of a single observation \mathbf{x} among all features (Štrumbelj and Kononenko, 2014). The Shapley value for the j -th feature is calculated by assessing its contribution to every possible feature combination $W \subseteq \{1, \dots, p\} \setminus \{j\}$ using the formula

$$\phi_j(\mathbf{x}) = \sum_{W \subseteq \{1, \dots, p\} \setminus \{j\}} \frac{|W|!(p - |W| - 1)!}{p!} (f_{W \cup \{j\}}(\mathbf{x}_{W \cup \{j\}}) - f_W(\mathbf{x}_W)),$$

where $f_W(\mathbf{x}_W) = E_{X_{-W}}[\hat{f}(\mathbf{x}_W, X_{-W})]$ is typically the PD function (Casalicchio et al., 2019). Since estimating the above equation is computationally expensive, various efficient approximation algorithms have emerged (Aas et al., 2021; Chen et al., 2023), including SHAP (Lundberg and Lee, 2017), which conceptualizes Shapley values as a solution to a quadratic program with equality constraints. To visualize the global feature effect of a feature on the prediction, Lundberg et al. (2020) proposed the SHAP dependence (SD) plot as an alternative to PD plots. SD plots visualize the feature values of the j -th feature on the x-axis and its corresponding SHAP values on the y-axis in a scatter plot, either for all or for a representative sample of observations (see Figure 4). Feature interactions between the feature of interest and other features influence the shape of the resulting point cloud. To highlight patterns in the point cloud that may indicate interactions with another feature, Lundberg et al. (2020) suggest color-coding the points by other (potentially interacting) features. However, identifying a clear trend or pattern in the color-coded point cloud becomes challenging in the presence of complex feature interactions (similar to ICE plots).

Shapley values can be quantified using either a marginal-based (similar to PDs) or a conditional-based approach. While the former might encounter extrapolation issues in the presence of feature correlations, the latter presents the challenge of estimating the conditional distribution. For a more thorough discussion on this topic, see Appendix J.

2.4 Quantification of Interaction Effects

Visualization methods for feature effects are usually limited to a maximum of two features. To understand which feature effects depend on other features due to interactions, we are interested in detecting and quantifying the strength of feature interactions. Building on the mean-centered PD function $\hat{f}_j^{PD,c}(\tilde{x}_j) = \hat{f}_j^{PD}(\tilde{x}_j) - \frac{1}{m} \sum_{k=1}^m \hat{f}_j^{PD}(\tilde{x}_j^{(k)})$ at each grid point $\tilde{x}_j \in \{\tilde{x}_j^{(1)}, \dots, \tilde{x}_j^{(m)}\}$, Friedman and Popescu (2008) introduced the H-Statistic as a global interaction measure that quantifies the interaction strength between a feature of interest \mathbf{x}_j and all other features \mathbf{x}_{-j} by

$$\hat{\mathcal{H}}_j^2 = \frac{\sum_{i=1}^n (\hat{f}^c(\tilde{\mathbf{x}}^{(i)}) - \hat{f}_j^{PD,c}(\tilde{x}_j^{(i)}) - \hat{f}_{-j}^{PD,c}(\tilde{\mathbf{x}}_{-j}^{(i)}))^2}{\sum_{i=1}^n (\hat{f}^c(\tilde{\mathbf{x}}^{(i)}))^2}. \quad (5)$$

Here, $\hat{f}^c(\tilde{\mathbf{x}}) = \hat{f}(\tilde{\mathbf{x}}) - \frac{1}{n} \sum_{i=1}^n \hat{f}(\tilde{\mathbf{x}}^{(i)})$ is the mean-centered prediction function and thus the denominator in Eq. (5) corresponds to the prediction function’s variance. Hence, $\hat{\mathcal{H}}_j^2$ quantifies how much of the prediction function’s variance can be attributed to the inter-

action between feature \mathbf{x}_j and all other features \mathbf{x}_{-j} .⁵ Thus, the H-Statistic can be used to rank features according to the global interaction strength. Despite its standardization between 0 and 1, it is unclear which value is generally considered a high or a low interaction. Furthermore, being derived from PD functions, the H-Statistic inherits the same disadvantages—particularly regarding extrapolation—as the PD plot itself.

3. Related Work

Below, we outline related work and define research gaps concerning our main contributions: GADGET (Section 4) for regional effects and PINT (Section 5) for testing interactions.

Related Work on Regional Effects. Regional explanations (Molnar et al., 2023; Britton, 2019; Hu et al., 2020; Scholbeck et al., 2024; Herbinger et al., 2022)—sometimes also referred to cohort explanations (Sokol and Flach, 2020)—explain subspaces of the entire feature space, thereby bridging the gap between local and global explanations. In general, we can distinguish between unsupervised and supervised methods to define these subspaces. Concerning the former, Britton (2019) suggests the use of an unsupervised clustering approach to group ICE curves based on their partial derivatives and apply a decision tree post hoc for explanation purposes. A similar approach has been introduced by Zhang et al. (2021). However, Herbinger et al. (2022) showed that such unsupervised approaches can produce misleading interpretations. To overcome these downsides, other works such as Molnar et al. (2023); Scholbeck et al. (2024); Herbinger et al. (2022) propose methods that are based on a supervised decision tree approach to define regions that are interpretable and distinct. The first two references focus on finding regions with few feature correlations and few non-linearities, respectively. Instead, the REPID method (Herbinger et al., 2022) minimizes feature interactions within each region and decomposes the global PD into regional PDs, ensuring the homogeneity of ICE curves within each region. However, REPID has its limitations: First, REPID generates feature-specific regions by minimizing interactions between a particular feature and all other features. Applying the method to a different feature of interest typically results in different regions, complicating interpretations when multiple features are of interest. Second, REPID relies on ICE and PD plots and thus inherits also their issues, such as extrapolation, potentially providing misleading interpretations if both feature interactions and correlations are present (Molnar et al., 2022). To overcome these limitations, we propose a more flexible framework in Section 4 that is capable of producing regions where interactions between multiple features are minimized and where different feature effect methods other than ICE and PD plots can be used.

Related Work on Detecting Feature Interactions. ANOVA is one of the most prominent statistical testing procedures to detect significant (two-way) interaction terms in statistical models. An algorithm to detect important feature interactions without a formal definition of a statistical test was introduced along with the standard functional ANOVA decomposition in Hooker (2004). Later, Mentch and Hooker (2017) introduced a statistical test for tree-based methods to detect additive structures and, thus, feature interactions. Henelius et al. (2016) introduced a statistical test for classification tasks to detect additive structures

5. If the feature \mathbf{x}_j does not interact with any other feature in \mathbf{x}_{-j} , the mean-centered prediction function can be decomposed as $\hat{f}^c(\tilde{\mathbf{x}}) = \hat{f}_j^{PD,c}(\tilde{\mathbf{x}}_j) + \hat{f}_{-j}^{PD,c}(\tilde{\mathbf{x}}_{-j})$ (see Section 2.2), leading to $\hat{\mathcal{H}}_j^2 = 0$.

between different subsets of features based on their earlier work Henelius et al. (2014). In their approach, feature values of the considered feature subset are permuted jointly within each class to break possible interactions with other feature subsets while preserving interactions within the subset. However, this approach is limited to classification tasks and breaks correlations with features outside of the regarded subset. Consequently, learned spurious interactions are detected by the algorithm as true feature interactions. Furthermore, statistical tests for two-way interaction detection have been introduced by Loh (2002) or Sorokina et al. (2008). While the former provides a formal statistical test statistic, the latter is based on a heuristic to detect two-way interactions by comparing the performance of a full and restricted model. To the best of our knowledge, a model-agnostic procedure designed for classification and regression tasks to detect interactions between a feature of interest \mathbf{x}_j and the remaining features \mathbf{x}_{-j} , as presented in Section 5, has not yet been introduced.

4. Generalized Additive Decomposition of Global Effects (GADGET)

Our proposed framework, GADGET, additively decomposes global feature effects by minimizing feature interactions to produce regional feature effects. One or more features can be used to derive interpretable subregions where aggregated regional feature effects better represent individual observations within each region. Below, we will introduce the GADGET methodology, applicable to various feature effect methods, by formally defining the required axiom and demonstrating its satisfaction by popular methods.

4.1 Measuring Interaction-Related Heterogeneity

Let $h(x_j, \mathbf{x}_{-j}^{(i)})$ be a function that quantifies the local effect of the j -th feature on the i -th observation at a specific value $x_j \in \mathcal{X}_j$, calculated by a function $h : \mathbb{R}^p \rightarrow \mathbb{R}$. We keep this function generic for now, but discuss in Sections 4.3 to 4.5 several specific choices, including the underlying local effect functions of PD, ALE, and SHAP dependence (SD) curves. For example, in the case of PD, h is defined by the i -th mean-centered ICE curve of the j -th feature at one specific grid point x_j (see Figure 2).

Let $\mathbb{E}[h(x_j, X_{-j})|\mathcal{A}_g]$ be the expected effect of the j -th feature at a specific value x_j regarding X_{-j} given the event $X \in \mathcal{A}_g$ where $\mathcal{A}_g = \iota_{g,1} \times \dots \times \iota_{g,p} \subseteq \mathcal{X}$ is a hypercube and the $\iota_{g,j} \subseteq \mathbb{R}$ are (potentially unbounded) intervals. Our goal is to find a partition of the feature space into subsets \mathcal{A}_g with little interaction-related heterogeneity. The deviation between the local feature effect and the expected feature effect at a specific value x_j for subspace \mathcal{A}_g is defined by a point-wise loss function, e.g., the squared distance:

$$\mathcal{L}_j(\mathcal{A}_g, x_j) = \sum_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g} \left(h(x_j, \mathbf{x}_{-j}^{(i)}) - \mathbb{E}[h(x_j, X_{-j})|\mathcal{A}_g] \right)^2. \quad (6)$$

Eq. (6) measures the heterogeneity of the local feature effects of feature \mathbf{x}_j at a specific value x_j by calculating the L2 loss as demonstrated in the left and middle plots of Figure 2. Here, the local feature effects are the mean-centered ICE curves and the respective mean-centered PD curve represents the expected effect of the feature of interest.

The risk function \mathcal{R}_j for the j -th feature and subspace \mathcal{A}_g is defined by aggregating the point-wise losses of Eq. (6) over a given m -dimensional grid of feature values $\tilde{\mathbf{x}}_j$:

$$\mathcal{R}_j(\mathcal{A}_g, \tilde{\mathbf{x}}_j) = \sum_{k:k \in \{1, \dots, m\} \wedge \tilde{\mathbf{x}}_j^{(k)} \in \iota_j} \mathcal{L}_j(\mathcal{A}_g, \tilde{\mathbf{x}}_j^{(k)}). \quad (7)$$

For instance, $\tilde{\mathbf{x}}_j$ can be m values sampled from the realizations of the j -th feature, an equidistant grid or quantiles, but we restrict the summation in Eq. (7) to values which fall inside the considered subspace \mathcal{A}_g , so $\iota_{g,j}$. It follows that Eq. (7) summarizes the heterogeneity of local feature effects of the feature \mathbf{x}_j for its entire feature range within the region \mathcal{A}_g in one number and therefore quantifies how much the feature's influence of individual observations varies from the expected (average) feature's influence on the predictions (e.g., right plot of Figure 2 in the case of mean-centered ICE and PD curves).

The loss in Eq. (6) measures the heterogeneity of the local effects, which can be decomposed into the main and interaction effects of involved features (see Eq. (1)). Since we are only interested in measuring the interaction-related heterogeneity of the feature of interest, disregarding variability caused by any other effects, the local feature effect function h needs to be defined in such a way that the heterogeneity of local feature effects measured by Eq. (6) and (7) are solely based on feature interactions between the feature of interest \mathbf{x}_j and the remaining feature in the data set. This is possible when Axiom 1 holds.⁶

Axiom 1 (Local Decomposability) *A local feature effect function $h : \mathbb{R}^p \rightarrow \mathbb{R}$ satisfies the local decomposability axiom if and only if $\forall i \in \{1, \dots, n\}$ the decomposition of the i -th local effect $h(x_j, \mathbf{x}_{-j}^{(i)})$ of feature \mathbf{x}_j at x_j solely depends on main and higher-order effects of feature \mathbf{x}_j :*

$$h(x_j, \mathbf{x}_{-j}^{(i)}) = g_j(x_j) + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}).$$

In Sections 4.3 to 4.5 we show that this axiom is fulfilled by the most popular feature effect methods, at least when properly centered. If Axiom 1 is satisfied, then Theorem 2 guarantees that the loss function defined in Eq. (6) quantifies the interaction-related heterogeneity of local effects (feature interactions) of feature \mathbf{x}_j at x_j within the regarded subspace \mathcal{A}_g . Specifically, Theorem 2 shows that the main effect $g_j(x_j)$ no longer affects the loss \mathcal{L}_j , rendering \mathcal{L}_j a pure measure of interaction-related heterogeneity.

Theorem 2 *If the local feature effect function $h(x_j, \mathbf{x}_{-j}^{(i)})$ satisfies Axiom 1, then the loss function $\mathcal{L}_j(\mathcal{A}_g, x_j)$ defined in Eq. (6) only depends on feature interactions between the feature \mathbf{x}_j at x_j and features in $-j$:*

$$\mathcal{L}_j(\mathcal{A}_g, x_j) = \sum_{i:\mathbf{x}^{(i)} \in \mathcal{A}_g} \left(\sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}) - \mathbb{E}[g_{W \cup j}(x_j, X_W) | \mathcal{A}_g] \right)^2.$$

The proof can be found in Appendix B.1.

6. We show in Sections 4.3 to 4.5 that common feature effect methods satisfy this general decomposition based on the functional ANOVA decomposition in Eq. (1).

Illustration. To better understand Axiom 1, Theorem 2, Eq. 6 and 7, we provide an illustrative example. Let $X_1, X_2, X_3 \sim \mathbb{U}(-1, 1)$ be independent, with the true underlying relationship defined by $Y = X_1 + X_2 + X_3 + X_1X_2 + \epsilon$, where $\epsilon \sim \mathbb{N}(0, 0.0025)$. We draw 20 observations and fit a correctly specified linear model to the data. We define the local feature effect function as the mean-centered ICE curves (Section 4.3) and use \mathbf{x}_1 as the feature of interest. The i -th ICE curve value at grid point \tilde{x}_1 is given by $\hat{f}(\tilde{x}_1, \mathbf{x}_{-1}^{(i)}) = \tilde{x}_1 + x_2^{(i)} + x_3^{(i)} + \tilde{x}_1x_2^{(i)}$, which corresponds to the decomposition in Eq. (1). Since this decomposition includes effects independent of feature \mathbf{x}_1 (i.e., $g_3(x_3^{(i)}) = x_3^{(i)}$), Axiom 1 does not hold. Therefore, the ICE curve needs to be centered by its mean: $\hat{f}^c(\tilde{x}_1, \mathbf{x}_{-1}^{(i)}) = \hat{f}(\tilde{x}_1, \mathbf{x}_{-1}^{(i)}) - \mathbb{E}[\hat{f}(X_1, \mathbf{x}_{-1}^{(i)})] = \tilde{x}_1 - \mathbb{E}[X_1] + (\tilde{x}_1 - \mathbb{E}[X_1])x_2^{(i)} \approx \tilde{x}_1 + \tilde{x}_1x_2^{(i)}$ which satisfies Axiom 1 with main effect $g_j(x_j) = \tilde{x}_1$ and the respective higher-order effect $g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}) = \tilde{x}_1x_2^{(i)}$. Figure 2 shows the mean-centered ICE curves and the mean-centered PD curve, which is calculated as the average over all mean-centered ICE curves. The loss function in Eq. (6) aggregates the squared distance between each mean-centered ICE curve and the PD curve at a specific grid point as demonstrated in the left and middle plot of Figure 2 for $\tilde{x}_1 = -1$. Since mean-centered ICE curves satisfy Axiom 1, this distance only measures interaction-related heterogeneity as stated in Theorem 2: $h(x_1, \mathbf{x}_{-1}^{(i)}) - \mathbb{E}[h(x_1, X_{-1})|\mathcal{X}] = \tilde{x}_1 + \tilde{x}_1x_2^{(i)} - \tilde{x}_1 - \tilde{x}_1\mathbb{E}[X_2] = \tilde{x}_1x_2^{(i)} - \tilde{x}_1\mathbb{E}[X_2] \approx \tilde{x}_1x_2^{(i)}$. Thus, the loss function and therefore also the risk function in Eq. (7), which aggregates over all valid grid points of the feature \mathbf{x}_j within the respective region (e.g., right plot Figure 2), solely depend on the interaction effect between the feature of interest (\mathbf{x}_1) and other features in the data set (here: \mathbf{x}_2).

4.2 The GADGET Algorithm

The goal of the GADGET algorithm is to find regions with little interaction-related heterogeneity. A user specifies two sets: $S \subseteq \{1, \dots, p\}$ containing the features of interest for which we want to receive representative regional feature effect plots by minimizing their interaction-related heterogeneity, and $Z \subseteq \{1, \dots, p\}$ containing features assumed to interact with those in S . Features in Z serve as split features when partitioning the feature space to minimize interactions with features in S . Features in $-S$ are usually assumed not to interact with other features and are thus only calculated and visualized on a global level. Algorithm 1 defines a single partitioning step of GADGET, which is inspired by the CART algorithm (Breiman et al., 1984) and is recursively repeated until a certain stop criterion is met (see Appendix D for more details). We greedily search for the best split point within the feature subset Z such that the interaction-related heterogeneity of feature subset S is minimized in the two resulting subspaces. The interaction-related heterogeneity is measured by the variance of the local feature effects of \mathbf{x}_j within the new subspace (see Eq. 7). Hence, for each candidate split feature z and candidate split point t , we sum up the risks of the two resulting subspaces for all features in S (line 6 in Algorithm 1) and then choose the split (\hat{t}, \hat{z}) that minimizes the interaction-related heterogeneity of all features $j \in S$ (line 9 in Algorithm 1). Note that Algorithm 1 is exemplarily defined for numeric features, however, GADGET can also handle categorical features which is described in detail in Appendix C.5.

6. S and Z can be distinct or partially or fully overlap with each other, see also Section 5.




Figure 2: Mean-centered ICE and PD curves of feature \mathbf{x}_1 of the described simulation example. Left: Illustrates the distance $d^{(1)}$ that is calculated within Eq. (6) between the local feature effect (here: $h(x_j, \mathbf{x}_{-j}^{(i)}) := \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)})$, i.e., the mean-centered ICE curve) and the expected effect within region \mathcal{A}_g (here: $\mathbb{E}[h(x_j, X_{-j})|\mathcal{X}] = \mathbb{E}[\hat{f}^c(x_j, X_{-j})|\mathcal{X}]$, i.e. the global mean-centered PD) at the first grid point $\tilde{x}_1 = -1$ for the first observation. Middle: The distances d are calculated for all mean-centered ICE curves at the first grid point and the squared values are summed up to obtain the loss function value for the first grid point as defined in Eq. (6). Right: The risk function value is calculated according to Eq. (7) by aggregating the loss function values over the valid grid points. This measures the heterogeneity of the mean-centered ICE curves of feature \mathbf{x}_1 , which quantifies the interaction-related heterogeneity between \mathbf{x}_1 and \mathbf{x}_{-1} based on Theorem 2.

Algorithm 1: Partitioning algorithm of GADGET

- 1: **input:** subspace $\mathcal{A} \subseteq \mathcal{X}$, risk \mathcal{R}_j , features S , splitting set Z , grids $\tilde{\mathbf{x}}_j \forall j \in S$
 - 2: **output:** subspaces $\mathcal{A}_l^{\hat{t}, \hat{z}}$ and $\mathcal{A}_r^{\hat{t}, \hat{z}}$
 - 3: **for** each feature indexed by $z \in Z$ **do**
 - 4: **for** every split t on feature \mathbf{x}_z **do**
 - 5: $\mathcal{A}_l^{t,z} = \{\mathbf{x} \in \mathcal{A} | \mathbf{x}_z \leq t\}$; $\mathcal{A}_r^{t,z} = \{\mathbf{x} \in \mathcal{A} | \mathbf{x}_z > t\}$
 - 6: $\mathcal{I}(t, z) = \sum_{j \in S} (\mathcal{R}_j(\mathcal{A}_l^{t,z}, \tilde{\mathbf{x}}_j) + \mathcal{R}_j(\mathcal{A}_r^{t,z}, \tilde{\mathbf{x}}_j))$
 - 7: **end for**
 - 8: **end for**
 - 9: Choose $(\hat{t}, \hat{z}) \in \arg \min_{t,z} \mathcal{I}(t, z)$
-

If the set of splitting features Z contains all features interacting with features in S , the GADGET algorithm eventually finds regions with no interaction-related heterogeneity (see Theorem 3). This implies that any feature interactions inherent in the feature effects of the features within S can be minimized, leaving only their main effects within each subspace. Thus, the joint regional feature effect $f_{S|\mathcal{A}_g}$ of all features in S within each subspace \mathcal{A}_g can be uniquely and additively decomposed into the respective univariate regional feature

effects $f_{j|\mathcal{A}_g}$ of all features in S :

$$f_{S|\mathcal{A}_g}(\mathbf{x}_S) = \sum_{j \in S} f_{j|\mathcal{A}_g}(\mathbf{x}_j). \quad (8)$$

While the global feature effects are estimated on the entire feature space \mathcal{X} , the regional feature effects are determined by conditioning on the respective subspace, i.e., on $X \in \mathcal{A}_g$.

Theorem 3 *Suppose the local feature effect function h satisfies Axiom 1. If the feature subset $-Z$ does not contain any features interacting with any features indexed by $j \in S$, i.e., $\forall j \in S$ and $\forall L \subseteq -Z$ with $-Z \subset \{1, \dots, p\} \setminus j$ and $\forall K \subseteq Z := \{\emptyset, 1, \dots, p\} \setminus \{-Z \cup j\}$ it holds: $g_{L \cup K \cup j}(X_j, X_L, X_K) = 0$, then after sufficiently many iterations, GADGET returns regions with $\mathcal{R}_j(\mathcal{A}_g, \tilde{\mathbf{x}}_j) = 0$ for every final region \mathcal{A}_g , i.e., the theoretical minimum of the objective function \mathcal{I} in Algorithm 1 is 0. The proof can be found in Appendix B.2.*

In practice, one rarely wants to run the algorithm to convergence. The question of how many partitioning steps should be performed depends on the underlying research question. If the user is more interested in reducing the interaction-related heterogeneity as much as possible, they might split rather deeply, depending on the complexity of interactions learned by the model. However, this might lead to many regions that are more challenging to interpret. If a small number of regions is of interest, one might prefer a shallow tree, thus reducing only the heterogeneity of the features that interact the most. More detailed recommendations on specifying these stop criteria are provided in Appendix D.

4.3 GADGET-PD

The PD plot is based on ICE curves (local feature effects). However, ICE curves do not satisfy Axiom 1, since the decomposition of the i -th ICE curve (i.e., $\hat{f}_j^{\mathbf{x}}(x_j) = \hat{f}(x_j, \mathbf{x}_{-j}^{(i)})$) also contains main or interaction effects of the i -th observation that are independent of feature \mathbf{x}_j (see Appendix B.3). These effects can be canceled out by centering ICE curves w.r.t. the mean of each curve, yielding the following local effect function in GADGET-PD:

$$h(x_j, \mathbf{x}_{-j}^{(i)}) = \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)}) = \hat{f}(x_j, \mathbf{x}_{-j}^{(i)}) - \mathbb{E}[\hat{f}(X_j, \mathbf{x}_{-j}^{(i)})|\mathcal{A}_g].$$

Theorem 4 *Mean-centered ICE curves satisfy the local decomposability axiom (Axiom 1). The proof is given in Appendix B.3.*

Note that the mean-centering constant depends on the region \mathcal{A}_g . In Figure 3, for example, we use feature \mathbf{x}_3 as a feature of interest in S and as a split feature in Z . Consequently, the visualized range of feature \mathbf{x}_3 is partitioned based on the split point identified by the GADGET algorithm and the mean-centering constant changed within the new subspace (i.e., $\mathbb{E}[\hat{f}(X_j, \mathbf{x}_{-j}^{(i)})|\mathcal{A}_g] \neq \mathbb{E}[\hat{f}(X_j, \mathbf{x}_{-j}^{(i)})]$), necessitating adjustment to avoid additive (non-interaction) effects in the new subspace.

The loss function used within GADGET-PD to minimize the interaction-related heterogeneity is defined by the variability of mean-centered ICE curves:

$$\mathcal{L}_j^{PD}(\mathcal{A}_g, x_j) = \sum_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g} \left(\hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)}) - \mathbb{E}[\hat{f}^c(x_j, X_{-j})|\mathcal{A}_g] \right)^2. \quad (9)$$

In Appendix B.4, we show that the REPID method discussed in Section 3 is a special case of GADGET-PD for the loss function in Eq. (9) with only one feature of interest (i.e., $S = j$) and where we consider all other features to be potential split features (i.e., $Z = -j$).

Illustration. To provide a better understanding of GADGET-PD, we consider a commonly used simulation example in a slightly modified form in the context of feature interactions (see e.g., Goldstein et al., 2015; Herbringer et al., 2022): Let $X_1, X_2, X_3 \sim \mathbb{U}(-1, 1)$ be independently distributed and the true underlying relationship be defined by $Y = 3X_1 \mathbb{1}_{X_3 > 0} - 3X_1 \mathbb{1}_{X_3 \leq 0} + X_3 + \epsilon$, with $\epsilon \sim \mathcal{N}(0, 0.09)$. We then draw 500 observations from these random variables and fit a feed-forward neural network (NN) with a single hidden layer of size 10 and weight decay of 0.001.⁷ The R^2 measured on a separately drawn test set of size 10000 following the same distribution is 0.94.

Figure 3 visualizes the global PD and ICE curves along with the regional ones after applying GADGET-PD. The global plots for features \mathbf{x}_1 and \mathbf{x}_3 show high heterogeneity in the underlying local effects, indicating that the global PD curve does not represent the respective ICE curves well. For example, the nearly horizontal global PD curve of \mathbf{x}_1 (grey line) indicates no influence on the model’s predictions. Consequently, it is not representative of the underlying local feature effects (ICE curves), which vary highly due to the feature interaction with \mathbf{x}_3 . We applied GADGET-PD with $S = Z = \{1, 2, 3\}$ to visualize the (regional) feature effect of all available features by considering all features as possible interaction (split) features. GADGET-PD performed one split with regard to $\mathbf{x}_3 \approx 0$. Thus, the correct split feature and its corresponding split point are found by GADGET-PD such that the interaction-related heterogeneity of all features in S is almost completely reduced and only the main effects within the subspaces remain. Hence, the joint mean-centered PD function $f_{S|\mathcal{A}_g}^{PD,c}$ within each final subspace \mathcal{A}_g can be decomposed into the respective 1-dimensional mean-centered PD functions as defined in Eq. (8). The decomposition is explained in more detail in Appendix C.1.

The main issue with local ICE curves and resulting global PD plots is the extrapolation problem when features are correlated. This problem and how it affects GADGET for different feature effect methods, will be explained in Section 4.6.

4.4 GADGET-ALE

While ALE curves avoid the issue of extrapolation seen in PD curves, they do not directly entail a local feature effect visualization that shows the heterogeneity induced by feature interactions, as seen in ICE curves for PD plots. However, ALE plots are based on derivatives of the prediction function regarding the feature of interest \mathbf{x}_j (see Eq. 3), which serve as local feature effects to capture the interaction-related heterogeneity. We use those derivatives within GADGET-ALE to receive more representative ALE curves in the final regions, yielding the following choice for the local feature effects: $h(x_j, \mathbf{x}_{-j}^{(i)}) = \frac{\partial f(x_j, \mathbf{x}_{-j}^{(i)})}{\partial x_j}$.

Theorem 5 *The local feature effects of GADGET-ALE satisfy the local decomposability axiom (Axiom 1). The proof is given in Appendix B.5.*

⁷ We tuned the size of the hidden layer over (3, 5, 10, 20, 30) and the weight decay value over (0.5, 0.1, 0.01, 0.001, 0.0001, 0.00001) via 5-fold cross-validation using grid search.




Figure 3: Visualization of applying GADGET with $S = Z = \{1, 2, 3\}$ to mean-centered ICE curves of the uncorrelated simulation example with $Y = 3X_1 \mathbb{1}_{X_3 > 0} - 3X_1 \mathbb{1}_{X_3 \leq 0} + X_3 + \epsilon$ where $\epsilon \sim \mathcal{N}(0, 0.09)$. Plots show mean-centered ICE and PD curves on the entire feature space (upper) and within regions after partitioning the feature space w.r.t. $\mathbf{x}_3 = -0.003$ (lower).

The derivatives in ALE are typically estimated by defining m intervals for feature \mathbf{x}_j and quantifying for each interval the prediction difference between the upper and lower boundary for observations lying in this interval (see Eq. 4). Hence, the loss function in Eq. (6) measures the variance of the derivatives of observations where the feature values of \mathbf{x}_j lie within the boundaries of the regarded interval and, thus, measures the interaction-related heterogeneity of local effects of feature \mathbf{x}_j within this interval. The risk function in Eq. (7) then aggregates this loss over all relevant intervals.

Equivalently to PD plots, ALE plots can be decomposed additively into main and interaction effects. Furthermore, if all feature interactions of features in S can be completely minimized (see Theorem 3), the joint mean-centered ALE function $f_{S|\mathcal{A}_g}^{ALE,c}$ within each final subspace \mathcal{A}_g can be decomposed into the 1-dimensional mean-centered ALE functions of features in S (see Eq. 8). More details on the decomposition including an illustrative example can be found in Appendix C.2.

4.5 GADGET-SD

While PD and ALE plots visualize the global feature effect for which we define the appropriate local feature effect function for the GADGET algorithm, the SD plot (Section 2.3) is comparable to an ICE plot in the sense that it shows all local feature effects instead of

an aggregated global feature effect. Here, we provide an estimate for the global effect and show the applicability of SD within GADGET which we term GADGET-SD.

Herren and Hahn (2022) showed that the Shapley value $\phi_j^{(i)}(x_j)$ of observation i for feature value x_j can be decomposed to

$$\phi_j^{(i)}(x_j) = g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j, |W|=k} g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}), \quad (10)$$

with $g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}) = \mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] + \sum_{k=0}^{|W|-1} \sum_{V \subset \{W \cup j\}, |V|=k} (-1)^{|W|-k} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}]$. These terms correspond to the pure feature interaction effects of feature \mathbf{x}_j as defined by the standard functional ANOVA decomposition (Hooker, 2004). For more details regarding the functional ANOVA representation of Shapley values, the interested reader is referred to Herren and Hahn (2022). This decomposition satisfies Axiom 1 and thus we choose $h = \phi_j^{(i)}$ for GADGET-SD.

Theorem 6 *The Shapley values satisfy the local decomposability axiom (Axiom 1). The proof can be found in Appendix B.6.*

It is important to note that the Shapley value is defined via expected values and thus its definition differs from those of ICE curves and derivatives for ALE, which are based on local predictions. Similarly to estimating the mean-centering constants for ICE curves, it follows that the expected values within the Shapley value estimation must consider the regarded subspace \mathcal{A}_g to acknowledge the full heterogeneity reduction due to interactions within each subspace. This means that we must recalculate the Shapley values after each partitioning step for each new subspace to receive regional SD effects in the final subspaces that are representative of the underlying main effects within each subspace. However, without recalculating the conditional expected values, we still minimize the unconditional expected value (i.e., the feature interactions on the entire feature space). The two different approaches lead to a different acknowledgment of interaction effects. In general, it can be said that the faster approach without recalculation will be less likely to detect feature interactions of higher order compared to the exact approach with recalculation. The differences of the two approaches are explained in Appendix C.3 and further discussed on a more complex simulation example in Section 6.2.

Based on Eq. (10), the global feature effect for the SD plot of feature \mathbf{x}_j at the feature value x_j can be defined by

$$f_j^{SD}(x_j) = \mathbb{E}_{X_W}[\phi_j(x_j)] = g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j, |W|=k} \mathbb{E}[g_{W \cup j}^c(x_j, X_W)]. \quad (11)$$

While there exist estimators for the global effect in PD and ALE plots, a pendant for the SD plot has not been introduced yet. Hence, a suitable estimator to estimate the expected value of Shapley values of feature \mathbf{x}_j in Eq. (11) must be chosen. Here, we use univariate GAMs with splines to estimate the expected value. Thus, the estimated GAM for feature \mathbf{x}_j represents the regional SD feature effect of feature \mathbf{x}_j within subspace \mathcal{A}_g (see Section 4.6).

The decomposition of Shapley values differs from that of mean-centered ICE curves in the way that feature interactions in Shapley values are fairly distributed between involved features, which leads to decreasing weights the higher the order of the interaction effect (see Eq. 10). In contrast, all interaction terms receive the same weight as the main effects in ICE curves. Hence, interactions of high order lead to less heterogeneity in SD plots compared to ICE plots. However, by using the marginal-based approach to calculate Shapley values, they can be decomposed by weighted PD functions (see Eq. 10 and Herren and Hahn, 2022). Hence, the same decomposition rules as defined for PD plots apply for the SD feature effect, as defined in Eq. (11). Meaning, the regional joint SD effect of features in S can be decomposed into 1-dimensional regional SD effect functions, as in Eq. (8) if all interaction-related heterogeneity is reduced. In this special case and if Shapley values are estimated by the marginal-based approach, it can be shown that $f_{j|\mathcal{A}_g}^{SD,c} = f_{j|\mathcal{A}_g}^{PD,c}$ (see Appendix C.3 for more details and an illustration in a simulated example).

4.6 Illustrative Comparison of Feature Effect Methods Within GADGET

We now illustrate the differences between the three discussed feature effect methods and how their underlying assumptions affect the resulting regional effects when GADGET is applied. We consider the simulation example from Section 4.3 with the true underlying relationship: $Y = 3X_1\mathbb{1}_{X_3>0} - 3X_1\mathbb{1}_{X_3\leq 0} + X_3 + \epsilon$. Here, all three features are jointly independent. The upper left plot in Figure 4 shows the ICE, global PD, and regional PD curves for feature \mathbf{x}_1 similar to Figure 3. The split found at $\mathbf{x}_3 \approx 0$ is optimal according to the true functional relationship of the data-generating process. The regional PD plots reflect the contradicting influence of feature \mathbf{x}_1 on the predictions due to the underlying feature interaction with \mathbf{x}_3 compared to the global PD plot. The respective split reduces the interaction-related heterogeneity almost completely (by 98%). Hence, the resulting regional marginal effects of \mathbf{x}_1 can be approximated well by the respective main effects (regional PD) and are more representative of the individual observations within each region.

One main disadvantage of ICE and PD plots is the extrapolation problem in the presence of feature interactions and correlations, as described in Section 2.3. To illustrate how this problem affects GADGET-PD, we consider the same simulation example as before, but with $X_1 = X_3 + \delta$ and $\delta \sim \mathcal{N}(0, 0.0625)$. We again draw 500 observations and train an NN with the same specification and receive an R^2 of 0.92 on a separately drawn test set of size 10000 following the same distribution. Thus, the high correlation between \mathbf{x}_1 and \mathbf{x}_3 barely influences the model performance. However, the upper right plot of Figure 4 clearly shows that ICE curves are extrapolating in low-density regions. The rug plot on the bottom indicates the distribution of feature \mathbf{x}_1 depending on feature \mathbf{x}_3 . Thus, the model was not trained on observations with small \mathbf{x}_1 values and simultaneously large \mathbf{x}_3 values, and vice versa. However, since we integrate over the marginal distribution, we also predict in these “out-of-distribution” areas. This leads to the so-called extrapolation problem and, hence, to uncertain predictions in extrapolated regions that must be interpreted with caution.

The four plots in the middle of Figure 4 illustrate the results for the same uncorrelated and correlated settings using GADGET-ALE. The upper plots show that the heterogeneity of the local effects (derivatives) before applying GADGET is very high, spanning across negative to positive values (grey boxplots) within each interval. With GADGET, we then

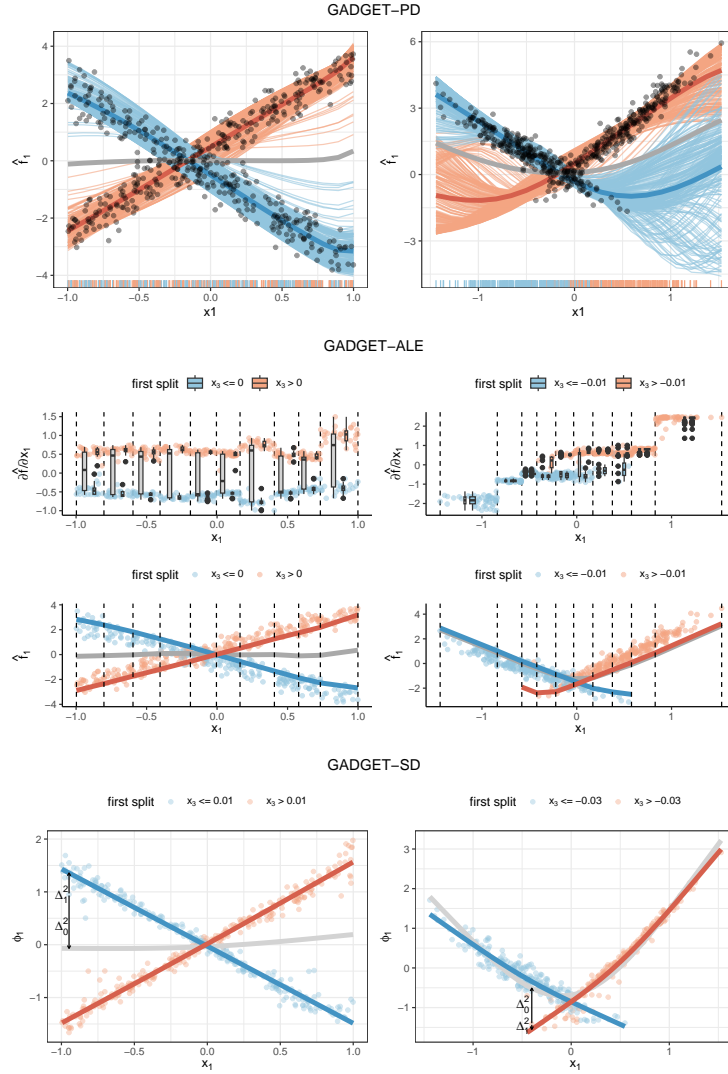


Figure 4: Local, global (grey), and regional effects for uncorrelated (left) and correlated (right) simulation settings. Local effects are colored w.r.t. the first split when GADGET is applied. The split feature is in all cases \mathbf{x}_3 . The blue color represents the left and orange the right region according to the split point. The thicker lines represent the respective regional PD curves. The rug plot and the black points in the upper plots show the distribution of \mathbf{x}_1 according to the split point and the underlying observational values, respectively. The two upper ALE plots visualize the heterogeneity of local feature effects (derivatives) while the two lower plots show the mean-centered global and regional ALE-curves.

partition the feature space w.r.t. one of the features in Z such that this interaction-related heterogeneity is minimized. GADGET-ALE finds the correct split feature and approxi-

mately the correct split point for both the uncorrelated and the correlated cases, and this split strongly reduces the heterogeneity of the local effects. While the shapes of the centered ALE curves look very similar to PD plots for the uncorrelated case, the ALE curves for the correlated case do not extrapolate and, thus, are more representative of the feature effect with regard to the underlying data distribution.

The lower plots in Figure 4 visualize the SD plots for feature \mathbf{x}_1 . Similarly to the least-square estimate in linear regression, we search for the GAM that minimizes the squared distance (see Δ^2 in two bottom plots in Figure 4) of the Shapley values of \mathbf{x}_1 . With GADGET, we now split such that the fitted GAMs within the two new subspaces minimize the squared distances between them and the Shapley values within the respective subspace. Since the GAMs are fitted on the Shapley values (local feature effects), in contrast to PDs, they do not extrapolate regarding \mathbf{x}_1 in the correlated scenario. However, Shapley values are based on expected values that must be estimated. If they are calculated using the marginal-based approach (as we do here), it is still possible that the predictions considered in the Shapley values extrapolate into sparse regions. Another difference to ICE curves that becomes visible in Figure 4 is the allocation of feature interactions. While the interaction effect of size 3 is fully attributed to both involved features for ICE curves, it is fairly distributed between \mathbf{x}_1 and \mathbf{x}_3 for Shapley values and thus the slopes of the regional effects for feature \mathbf{x}_1 are halved compared to the regional PD curves. In which way this difference affects the GADGET algorithm will be analyzed in Section 6.2.

Definitions and descriptions of the mean and interaction-related heterogeneity estimates and visualization techniques for each feature effect method can be found in Appendix C.4.

4.7 Quantifying Feature Interactions

In addition to visualizing regional effects, the GADGET algorithm quantifies feature interactions, for which we introduce several measures (inspired by Herbinger et al., 2022). We quantify how much interaction-related heterogeneity has been reduced in a single split of GADGET for a considered feature $j \in S$,

$$I(\mathcal{A}_P, \tilde{\mathbf{x}}_j) = \frac{\mathcal{R}_j(\mathcal{A}_P, \tilde{\mathbf{x}}_j) - \mathcal{R}_j(\mathcal{A}_l, \tilde{\mathbf{x}}_j) - \mathcal{R}_j(\mathcal{A}_r, \tilde{\mathbf{x}}_j)}{\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j)},$$

where P is the parent node and l and r are its left and right children, respectively. We express risk reduction via the split relative to the risk on the entire feature space (root node). Instead of considering only a single partitioning step, one might be more interested in how much interaction-related heterogeneity reduction a specific split feature $z \in Z$ is responsible for w.r.t. all performed partitioning steps of an entire tree. For a given splitting feature z we now simply sum up the reduction terms for all splits in which z occurs, which we denote by $\mathcal{B}_z \subset \{\mathcal{A}_1, \dots, \mathcal{A}_G\}$, where G is the number of splitting nodes in the tree:

$$I_{z,j}(\tilde{\mathbf{x}}_j) = \sum_{\mathcal{A}_P \in \mathcal{B}_z} I(\mathcal{A}_P, \tilde{\mathbf{x}}_j).$$

We obtain the overall interaction-related heterogeneity reduction I_z of the z -th split feature by relating the risk reduction of all its splits over all considered features in S to the

total risk of the root node:

$$I_z = \frac{\sum_{\mathcal{A}_P \in \mathcal{B}_z} \sum_{j \in S} (\mathcal{R}_j(\mathcal{A}_P, \tilde{\mathbf{x}}_j) - \mathcal{R}_j(\mathcal{A}_l, \tilde{\mathbf{x}}_j) - \mathcal{R}_j(\mathcal{A}_r, \tilde{\mathbf{x}}_j))}{\sum_{j \in S} \mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j)}.$$

For example, in our previous example in Figure 3, the interaction-related heterogeneity reduction of the first split ($I(\mathcal{X}, \tilde{\mathbf{x}}_1)$) for feature \mathbf{x}_1 is 0.986, which means that almost all of the interaction-related heterogeneity of \mathbf{x}_1 is reduced after the first split. Furthermore, we obtain $I_{3,1}(\tilde{\mathbf{x}}_1) = I(\mathcal{X}, \tilde{\mathbf{x}}_1) = 0.986$, since $z = 3$ was only used once for splitting, while the interaction-related heterogeneity reduction of all three features is $I_3 = 0.99$ for $z = 3$.

Goodness of Fit. A further aggregation level would be to sum up $I_{z,j}(\tilde{\mathbf{x}}_j)$ over all $z \in Z$ and, thus, obtain the interaction-related heterogeneity reduction for feature \mathbf{x}_j between the entire feature space and the terminal nodes. This is related to the concept of R^2 , which is a well-known measure in statistics to quantify the goodness of fit. We apply this concept here to quantify how well the terminal regional effect curves (in the terminal subspaces $\mathcal{B}_t \subset \{\mathcal{A}_1, \dots, \mathcal{A}_G\}$) fit the underlying local effects compared to the global feature effect curve on the entire feature space. We distinguish between the feature-related R_j^2 , which represents the goodness of fit for the feature effects of feature \mathbf{x}_j :

$$R_j^2 = \sum_{z \in Z} I_{z,j}(\tilde{\mathbf{x}}_j) = 1 - \frac{\sum_{\mathcal{A}_t \in \mathcal{B}_t} \mathcal{R}_j(\mathcal{A}_t, \tilde{\mathbf{x}}_j)}{\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j)},$$

and the R_{Tot}^2 , which quantifies the goodness of fit for the feature effects of all features in S :

$$R_{Tot}^2 = \sum_{z \in Z} I_z = 1 - \frac{\sum_{\mathcal{A}_t \in \mathcal{B}_t} \sum_{j \in S} \mathcal{R}_j(\mathcal{A}_t, \tilde{\mathbf{x}}_j)}{\sum_{j \in S} \mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j)}.$$

The right-hand versions of both formulas follow from the cancellation of parent vs. children terms (except for root and terminal nodes) when the left-hand version is summed across the whole tree. Both R^2 measures take values between 0 and 1, with values close to 1 signaling that almost all heterogeneity in the final subspaces compared to the entire feature space has been reduced—either for a specific feature of interest (R_j^2) or for all features of interest (R_{Tot}^2). In our example, R_1^2 for feature \mathbf{x}_1 is the same as for $I_{3,1}(\mathbf{x}_1)$, since GADGET performed only one split. The total interaction-related heterogeneity reduction over all features in S is $R_{Tot}^2 = I_3 = 0.99$. Hence, the interaction-related heterogeneity of all features in S has been reduced by 99% after the first split.

5. A Procedure to Detect Global Feature Interactions

As shown in Section 4.7, GADGET quantifies feature interactions between the feature subsets S and Z . Choosing the features to include in S and Z depends on the underlying research question. If the user is interested in how a specific set of features (S) influences the model’s predictions depending on another user-defined feature set (Z), then S and Z are chosen based on domain knowledge. However, the user may not know which relationships the ML model has learned. Thus, choosing S and Z based on domain knowledge does not

ensure consideration of all interacting features. If our goal is to minimize feature interactions between all features to potentially additively decompose the prediction function into univariate effects, we can define $S = Z = \{1, \dots, p\}$. With that choice, we aim to reduce the overall interaction-related heterogeneity in all features (S), since we also consider all features to be possible interacting features (Z). For high-dimensional settings, this has the disadvantage of being computationally expensive and we might “detect” spurious interactions and therefore produce less stable regional splits. Hence, we must define beforehand the subset of features that actually interact. Given that features typically interact with each other and all involved features will usually show heterogeneity in their local effects while being responsible for the heterogeneity of other involved features, we will choose $S = Z$. Thus, we will only use S as the globally interacting feature subset to be defined.

5.1 The PINT Procedure

We introduce a new permutation-based interaction detection (PINT) procedure to define the interacting feature subset S . PINT is based on the idea of a non-parametric permutation test (Ernst, 2004). It can be applied to any feature effect method satisfying Axiom 1.

With the risk function defined in Eq. (7) and based on the chosen local feature effect method, we can quantify the interaction-related heterogeneity of each feature \mathbf{x}_j within the feature space \mathcal{X} . Due to correlations between features, the estimated heterogeneity might also include spurious interactions. Thus, we must define a null distribution to determine which heterogeneity is actually due to feature interactions (i.e., significant w.r.t. the null distribution) and which heterogeneity is due to other reasons, such as correlations or noise. Thus, the null distribution needs to retain the underlying properties of the data.

Generally, we do not know the true functional relationship and there are infinitely many distributions compliant with a null hypothesis of no interaction-related heterogeneity: $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j) = 0$. We use the prediction model and a permutation-based approach to approximate the variability of the risk \mathcal{R}_j under this null hypothesis. We follow a similar approach to Altmann et al. (2010) who suggested a permutation-based null distribution to detect important features. To obtain the null distribution, the target variable y is permuted (line 4 in Algorithm 2), which breaks the association between the target and the features while keeping the feature distribution intact. Hence, no more interactions are present between \mathbf{x}_j and \mathbf{x}_{-j} , but the association between the features remains. Thus, when performing a refit on the data set \tilde{D} with the permuted target variable, potential heterogeneity of local effects is only due to randomly learned effects of the prediction model depending on the feature distribution and noise and not caused by true feature interactions. Therefore, calculating $\tilde{\mathcal{R}}_j$ based on \tilde{D} measures the risk of \mathbf{x}_j under the null (lines 5-8 in Algorithm 2). To obtain the null distribution of risk values for \mathbf{x}_j , this procedure is repeated s times.⁸

Since we are interested in defining the feature subset S of all interacting features, we calculate $\tilde{\mathcal{R}}_j$ for all features $j \in \{1, \dots, p\}$. Since all feature interactions between y and \mathbf{x} are

8. An alternative approach would be to randomly permute \mathbf{x}_j instead of y to break the interaction effects between \mathbf{x}_j and \mathbf{x}_{-j} . However, this approach is computationally more expensive and does not only break the respective interaction effects but also the correlations between \mathbf{x}_j and \mathbf{x}_{-j} , and thus spurious interactions will not be detected. Using conditional samples (Molnar et al., 2023) or knock-offs (Watson and Wright, 2021) instead of random permutations solves the second problem. However, this approach is even more computationally expensive and becomes infeasible with a higher number of features.

Algorithm 2: PINT

```

1: input: data set  $\mathcal{D}$ , prediction function  $\hat{f}$ , number of permutations  $s$ ,
   risk function  $\mathcal{R}_j$ , significance level  $\alpha$ 
2: output: feature subset  $S$ 
3: for  $k \in \{1, \dots, s\}$  do
4:   permute target  $y$  of  $\mathcal{D}$  and obtain permuted  $\tilde{y}^k$  and  $\tilde{\mathcal{D}}^k$ 
5:   refit model on  $\tilde{\mathcal{D}}^k$  to obtain the prediction function  $\tilde{f}_{\tilde{\mathcal{D}}^k}^k$ 
6:   for  $j \in \{1, \dots, p\}$  do
7:     calculate risk  $\tilde{\mathcal{R}}_j^k = \mathcal{R}_j^k(\mathcal{X}, \tilde{\mathbf{x}}_j, \tilde{f}_{\tilde{\mathcal{D}}^k}^k)$  for the  $j$ -the feature achieved by  $\tilde{f}_{\tilde{\mathcal{D}}^k}^k$ 
8:   end for
9: end for
10: for  $j \in \{1, \dots, p\}$  do
11:   calculate risk  $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}})$  for  $j$ -the feature based on  $\mathcal{D}$  and  $\hat{f}$ 
12:   sort (increasing) all  $\tilde{\mathcal{R}}_j^k$ , set  $(1 - \alpha)$ -quantile  $z_j^{1-\alpha} = \tilde{\mathcal{R}}_j^{[(s+1)(1-\alpha)]}$ 
13:   add  $j$  to  $S$ , if and only if  $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}}) > z_j^{1-\alpha}$ 
14: end for
    
```

broken by permuting y , the null distribution of risk values can be estimated for all features by the described s permuting and refitting steps. To determine the interacting feature subset S , the respective risk value $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}})$, is calculated based on the unpermuted data set \mathcal{D} and the originally fitted model on this data set. The value of $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}})$ is then compared to the $(1 - \alpha)$ -quantile, i.e., $z_j^{1-\alpha}$, of the non-parametric null distribution. If the risk value of \mathbf{x}_j exceeds the respective quantile, then feature interactions between \mathbf{x}_j and \mathbf{x}_{-j} are considered to be highly relevant (lines 11-13 in Algorithm 2).

5.2 Theoretical Properties

We emphasize that PINT is a detection *procedure*, not a formal statistical test. This distinction is necessary because the permutation approach generates data from only one of the many possible distributions compliant with the null hypothesis. We may nevertheless ask how the procedure's type I and type II errors behave. In the following, we give conditions under which PINT is valid as a statistical test. An empirical validation is given in Section 7 and Appendix F.

A difficulty in formally analyzing the procedure is that it depends on the unspecified ML algorithm that induces \hat{f} . We can emphasize this by writing $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}_n})$ for the risk for feature value \mathbf{x}_j achieved by the algorithm \hat{f} run on a data set⁹ $\mathcal{D}_n = \{(X^{(i)}, Y^{(i)})\}_{i=1}^n$, generated i.i.d. from $\mathbb{P}_{X,Y}$. Define $f(\mathbf{x}) = \mathbb{E}[Y \mid X = \mathbf{x}]$ and $\mathcal{D}_n^\Pi = \{(X^{(i)}, Y_{\Pi(i)}^{(i)})\}_{i=1}^n$ where Π is a random permutation of $\{1, \dots, n\}$ independent of \mathcal{D}_n . We can now impose abstract, but interpretable conditions on the ML algorithm under which the test is valid.

9. With a slight abuse of notation, we use \mathcal{D}_n to define the data set as a random variable.

Theorem 7 *Suppose $\mathbb{P}_{X,Y}$ is a probability measure where $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, f) = 0$. The type I error probability of PINT is upper bounded by $\alpha + \epsilon$, where*

$$\epsilon = \max_{t>0} \left[\mathbb{P} \left(\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}_n}) > t \right) - \mathbb{P} \left(\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}_n^\Pi}) > t \right) \right].$$

The proof can be found in Appendix E.

Intuitively, \mathcal{D}_n^Π is a data set with no feature effects. The extra term ϵ is zero if the algorithm \hat{f} is less likely to find a spurious interaction in unpermuted data compared to a case where no effects are present. Often, this is natural: most learners detect less of the spurious effect when there are other effects to focus on. In this case, the PINT procedure is a valid statistical test. In fact, we conjecture that the test is conservative (i.e., type I error is strictly smaller than α) in most cases, which is in line with our empirical observations (see Section 7 and Appendix F). However, ϵ may be strictly positive when features are strongly dependent, and the spurious effect is easier to learn than the true one. An example of such a case is given in Appendix F.1.

Theorem 8 *Suppose $\mathbb{P}_{X,Y}$ is a probability measure where $\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, f) \neq 0$. If the algorithm \hat{f} is such that for every $\epsilon > 0$*

$$\mathbb{P} \left(\sup_x |\hat{f}_{\mathcal{D}_n^\Pi}(\mathbf{x}) - \mathbb{E}[Y]| > \epsilon \right) \rightarrow 0 \tag{12}$$

and there exists $c > 0$ such that

$$\mathbb{P} \left(\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j, \hat{f}_{\mathcal{D}_n}) > c \right) \rightarrow 1, \tag{13}$$

the type II error probability tends to 0. The proof can be found in Appendix E.

Recall that \mathcal{D}_n^Π is a data set with no feature effects, in which case we would have $f(\mathbf{x}) = \mathbb{E}[Y]$ for all \mathbf{x} . Condition (12) essentially requires the inducing algorithm of \hat{f} to perfectly learn constant functions. When there are feature effects, the algorithm does not have to be perfect, though. Condition (13) merely requires it to capture at least some of the interaction effect between x_j and other features.

6. Empirical Validation and Analyses of GADGET

We now investigate different hypotheses to (1) empirically validate that GADGET generally minimizes feature interactions and (2) show how different characteristics of the underlying data—like feature correlations and true functional relationships—might influence GADGET, depending on the chosen feature effect method. The structure of the following sections and the concrete definition of the hypotheses is based on (2). However, the simulation examples themselves are designed in such a way that we know the underlying ground-truth of feature interactions for each example. Thus, we are able to empirically validate that GADGET generally minimizes feature interactions.

Definition of Stop Criteria for GADGET. In all our experiments and real-world applications, we control the number of partitioning steps performed by GADGET based on (1) the tree depth and the minimum number of observations per leaf node which are hyperparameters typically used by decision trees, and (2) an early stopping mechanism where a further split is only performed if the relative improvement of the split to be performed is at least $\gamma \in [0, 1]$ times the total relative interaction-related heterogeneity reduction of the previous split: $\gamma \times \frac{\sum_{j \in S} (\mathcal{R}_j(A_P, \mathbf{x}_j)) - \mathcal{I}(\hat{t}, \hat{z})}{\sum_{j \in S} (\mathcal{R}_j(\mathcal{X}, \mathbf{x}_j))}$. The higher we choose γ , the fewer partitioning steps will be performed. See Appendix D for more details and recommendations.

6.1 Influence of Correlated Features

Hypothesis. With increasing correlation between features, results become less stable due to extrapolation for methods using the marginal distribution for integration (especially PD, as shown in Section 4.6, but also SD) compared to methods using the conditional distribution (such as ALE). In this context, “stability” denotes the ability of the different GADGET variants to find the correct split feature and split point over various repetitions.

Experimental Setting. We use the (simple) simulation example of Section 4 for four different correlation coefficients ρ_{13} between X_1 and X_3 : 0, 0.4, 0.7, and 0.9. The data is generated as follows: Let $X_2, X_3 \sim \mathcal{U}(-1, 1)$ be independently distributed and $X_1 = c \cdot X_3 + (1 - c) \cdot Z$ with $Z \sim \mathcal{U}(-1, 1)$, where c is chosen between 0 and 0.7, to achieve the ρ_{13} values. The true underlying relationship is defined as before by $Y = 3X_1 \mathbb{1}_{X_3 > 0} - 3X_1 \mathbb{1}_{X_3 \leq 0} + X_3 + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.09)$. We draw 1000 observations and fit a GAM with correctly specified main and interaction effects and an NN with the previously defined specifications. We repeat the experiment 30 times. We apply GADGET to each setting and model within each repetition using PD, ALE, and SD as feature effect methods and set $S = Z = \{1, 2, 3\}$. As stopping criteria, we choose a maximum tree depth of 6, a minimum number of observations per leaf of 40, and set the improvement parameter γ to 0.2.

Results. Figure 5 shows that independent of the model or correlation degree, \mathbf{x}_2 has (correctly) never been considered as split feature. For correlation strengths ρ_{13} between 0 and 0.7, \mathbf{x}_3 is always chosen as the only split feature, with I_3 taking values between 0.75 and 1 and thus reducing most of the interaction-related heterogeneity of all features with one split. Thus, for low to medium correlations, there are only minor differences between the various feature effect methods and models. However, the observed behavior changes substantially for $\rho_{13} = 0.9$. For the correctly specified GAM, we still receive fairly consistent results, apart from one repetition where \mathbf{x}_1 is chosen as the split feature when PD is used. In contrast, there is more variation of I_3 for the NN. For SD, \mathbf{x}_1 is chosen once for splitting, and for PD, this is also the case for 30% of all repetitions (see Table 1). Using ALE, GADGET always correctly performs one split with respect to \mathbf{x}_3 . Additionally, GADGET performs a second split once when PD is used and 7 times when SD is used. Thus, for high correlations between features, we already observe in this very simple setting that the extrapolation problem influences the splitting within GADGET for effect methods based on marginal distributions, while methods based on conditional distributions such as ALE are less affected. This is particularly relevant for learners that model very locally (e.g. NNs) and, thus, tend to have wiggly prediction functions and oscillate in extrapolating regions.




Figure 5: Boxplots showing the interaction-related heterogeneity reduction I_z per split feature over 30 repetitions when PD, ALE, or SD is used in GADGET. Columns refer to correlation ρ_{13} , rows refer to fitted ML model.

This is also notable in the split value range, which increases for the NN in the case of high correlations for PD and SD but not for ALE (see Table 1).

		Split feature \mathbf{x}_3 in %			Split value range			MSE
Model	ρ_{13}	ALE	PD	SD	ALE	PD	SD	mean (sd)
GAM	0	1.00	1.00	1.00	0.10	0.10	0.12	0.329 (0.008)
GAM	0.4	1.00	1.00	1.00	0.13	0.10	0.10	0.201 (0.003)
GAM	0.7	1.00	1.00	1.00	0.09	0.10	0.09	0.137 (0.002)
GAM	0.9	1.00	0.97	1.00	0.11	0.11	0.12	0.107 (0.001)
NN	0	1.00	1.00	1.00	0.10	0.09	0.09	0.152 (0.018)
NN	0.4	1.00	1.00	1.00	0.08	0.08	0.07	0.124 (0.007)
NN	0.7	1.00	1.00	1.00	0.11	0.10	0.10	0.111 (0.005)
NN	0.9	1.00	0.70	0.97	0.10	0.19	0.15	0.104 (0.003)

Table 1: Overview of the frequency of \mathbf{x}_3 as the first split feature over 30 repetitions, along with the corresponding split value range. The last column presents the model’s test performance as the mean (standard deviation) of the mean squared error (MSE).

6.2 Influence of Higher-Order Interaction Effects

Hypothesis. While SD puts less weight on interactions with increasing order, all interactions (regardless of the order) receive the same weight in PD and ALE. Hence, when using GADGET-SD, we may not be able to detect high order interactions—especially when using the approximation without recalculation after each split (see Section 4.5 and Appendix C.3). However, it should be more likely to detect these interactions with PD and ALE.

Experimental Setting. To investigate this hypothesis, we consider 5 independent features with $X_1 \sim \mathbb{U}(0, 1)$ and $X_2, X_3, X_4, X_5 \sim \mathbb{U}(-1, 1)$ and draw 1000 samples. The true functional relationship of the data-generating process is defined by a series of interactions between different features: $y = f(\mathbf{x}) + \epsilon$ with $f(\mathbf{x}) = \mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 \leq 0} \mathbb{1}_{\mathbf{x}_4 > 0} + 4\mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 \leq 0} \mathbb{1}_{\mathbf{x}_4 \leq 0} - \mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 > 0} \mathbb{1}_{\mathbf{x}_5 \leq 0} \mathbb{1}_{\mathbf{x}_2 > 0} - 3\mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 > 0} \mathbb{1}_{\mathbf{x}_5 \leq 0} \mathbb{1}_{\mathbf{x}_2 \leq 0} - 5\mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 > 0} \mathbb{1}_{\mathbf{x}_5 > 0}$ and $\epsilon \sim \mathcal{N}(0, 0.01 \cdot \text{var}(f(\mathbf{x})))$. These interactions can be seen as one hierarchical structure between all features, where the slope of \mathbf{x}_1 depends on the subspace defined by the interacting features (see Figure 6).

We fit an XGBoost (XGB) model with correctly specified feature interactions and a random forest (RF) on the data set. Subsequently, we apply GADGET to each model using PD, ALE, SD with, and SD without recalculation after each split. In 30 repetitions of the experiment, XGB achieved an MSE of 0.068 (0.009 standard deviation) and the RF 0.121 (0.012 standard deviation) on a separate test set with the same distribution. For GADGET, we consider one feature of interest $S = 1$ and all other features as potential interacting features $Z = \{2, 3, 4, 5\}$. As stop criteria, we set the maximum tree depth to 7, the minimum number of observations to 40, and $\gamma = 0.1$. If the underlying model learned the effects of the true functional relationship of the data-generating process correctly, GADGET should split as shown in Figure 6 to minimize the interaction-related heterogeneity of \mathbf{x}_1 .

Results. All methods used \mathbf{x}_3 as the first split feature in all repetitions. Table 2 shows that a second-level split was performed in only 10% of the repetitions when GADGET-SD without recalculation is applied on the XGB model, compared to about 90% for all other methods. A similar pattern appears with the RF model, but with more variation in the frequencies, possibly due to different learned effects. If a second-level split is performed, all methods always choose the correct split features \mathbf{x}_4 and \mathbf{x}_5 , indicating that GADGET generally minimizes feature interactions (see Figure 6). Table 3 shows similar differences in the third-level split frequencies between SD without recalculation and other methods. Table 4 confirms that SD without recalculation shows high variation in final subspaces, with a median of 2, indicating it stops after the first split (with \mathbf{x}_3). Other GADGET variants, like PD and SD with recalculation, typically reach the correct number of subspaces (5), while ALE tends to split slightly deeper.

To summarize, when SD without recalculation is used, the two-way interaction with \mathbf{x}_3 is primarily detected, while features of a third-order (\mathbf{x}_4 and \mathbf{x}_5) or fourth-order (\mathbf{x}_2) interaction are rarely considered for splitting. In contrast, the other three methods frequently detect higher-order interactions. This supports our hypothesis regarding higher-order effects and the theoretical differences among the feature effect methods. Note that recalculating the Shapley values after each split reduces the order of interactions. For example, recalculating Shapley values in the subspace $\{\mathcal{X} | \mathbf{x}_3 \leq 0\}$ reduces the three-way interaction $\mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_3 \leq 0} \mathbb{1}_{\mathbf{x}_4 > 0}$ to the two-way interaction $\mathbf{x}_1 \cdot \mathbb{1}_{\mathbf{x}_4 > 0}$ and thus the weight of the interaction increases for the next split. As a result, the outcomes are similar to those for PD. However, recalculation can be computationally expensive.

6.3 Influence of Spurious Interactions

Hypothesis. When using PINT to pre-select $S = Z \subseteq \{1, \dots, p\}$, we are more likely to actually split according to feature interactions compared to when choosing all features




Figure 6: Explanation of the true functional relationship of the data-generating process. The green squares represent the 5 final subspaces which contain linear effects of feature \mathbf{x}_1 .

Model	Method	$\mathcal{A}_l^{t,z}$		$\mathcal{A}_r^{t,z}$	
		z	Freq.	z	Freq.
XGB	ALE	x4	0.93	x5	0.93
XGB	PD	x4	0.90	x5	0.90
XGB	SD_not_rc	x4	0.10	x5	0.10
XGB	SD_rc	x4	0.87	x5	0.87
RF	ALE	x4	0.80	x5	0.90
RF	PD	x4	0.63	x5	0.73
RF	SD_not_rc	x4	0.03	x5	0.13
RF	SD_rc	x4	0.73	x5	0.90

Table 2: Frequencies of splits relative to root node by split feature z over 30 repetitions in the left $\mathcal{A}_l^{t,z}$ and right $\mathcal{A}_r^{t,z}$ subspaces after the first split when GADGET with ALE, PD, SD with recalculation (rc) and without recalculation (not_rc) is used.

Model	Method	z	Rel. Freq.
XGB	ALE	x2	0.93
XGB	PD	x2	0.90
XGB	SD_not_rc	x2	0.10
XGB	SD_rc	x2	0.87
RF	ALE	x2	0.83
RF	ALE	x4	0.03
RF	PD	x2	0.67
RF	SD_not_rc	x2	0.10
RF	SD_rc	x2	0.83

Table 3: Frequencies of splits relative to root node by split feature z over all 30 repetitions in the subspace $\{\mathcal{X} | \mathbf{x}_3 > 0 \cap \mathbf{x}_5 \leq 0\}$ on third tree depth when GADGET is applied with ALE, PD, SD_not_rc and SD_rc.

		No. of Subspaces		
Model	Method	Min	Max	Med
XGB	ALE	3	15	7
XGB	PD	2	7	5
XGB	SD_not_rc	2	7	2
XGB	SD_rc	2	8	5
RF	ALE	3	16	9
RF	PD	2	11	5
RF	SD_not_rc	2	11	2
RF	SD_rc	3	12	5

Table 4: Minimum, maximum and median number of final subspaces over 30 repetitions after GADGET is applied with ALE, PD, SD_not_rc and SD_rc.

($S = Z = \{1, \dots, p\}$) or using the values of the H-Statistic—which was defined in Section 2.4—for pre-selection, especially when potential spurious interactions are present.

Experimental Setting. We reuse the simulation example described in Section 5: Consider four features with $X_1, X_2, X_4 \sim U(-1, 1)$ and $X_3 = X_2 + \epsilon$ with $\epsilon \sim N(0, 0.09)$. For 30 repetitions, we draw $n = \{300, 500\}$ observations and create the dependent variable,

including a potential spurious interaction between \mathbf{x}_1 and \mathbf{x}_3 (as \mathbf{x}_1 and \mathbf{x}_2 interact, but \mathbf{x}_2 and \mathbf{x}_3 are correlated)¹⁰: $\mathbf{y} = \mathbf{x}_1 + \mathbf{x}_2 + \mathbf{x}_3 - 2\mathbf{x}_1\mathbf{x}_2$. For each sample size n and each repetition, we fit an SVM with an RBF kernel, cost parameter $C = 1$, and choose the inverse kernel width based on the data. The performance of the model, measured by the mean (standard deviation) of the MSE on a test set of size 100000 across all repetitions, is 0.028 (0.010) for $n = 300$ and 0.027 (0.010) for $n = 500$. We calculate PINT using PD, ALE, and SD for each repetition and sample size with $s = 100$ and $\alpha = 0.05$ by approximating the null distribution as described in Section 5. We apply GADGET using PD, ALE, and SD with recalculation, where $S = Z$ is based on the feature subset chosen by PINT for each method. We compare these results by considering all features as features of interest and potential split features (i.e. $S = Z = \{1, \dots, p\}$). We use a maximum tree depth of 6, minimum number of observations of 40, and $\gamma = 0.15$ as stop criteria.

Results. The two left plots in Figure 7 show that \mathbf{x}_3 and \mathbf{x}_4 are always correctly identified as insignificant (according to the chosen α level), while the interacting features \mathbf{x}_1 and \mathbf{x}_2 are always significant and considered in S (apart from a few exceptions for ALE). The sample size does not seem to have a clear influence on PINT in this setting. The right plots in Figure 7 show that even the H-Statistic values of the non-influential and uncorrelated feature \mathbf{x}_4 are larger than 0 for both sample sizes. The H-Statistic values of \mathbf{x}_3 could support its inclusion in S , depending on the chosen threshold. Since this choice is not clear for the H-Statistic, it is not very suitable as a pre-selection method for GADGET.

Figure 8 illustrates that we correctly only consider \mathbf{x}_1 and \mathbf{x}_2 when PINT is applied upfront, while GADGET also splits w.r.t. \mathbf{x}_3 for all settings and (in some cases) even w.r.t. \mathbf{x}_4 if PINT is not applied upfront. The influence of these two features appears to be stronger for smaller sample sizes, as indicated by I_2 . Among the various effect methods used in GADGET, PD and SD attribute most of the reduction in heterogeneity to \mathbf{x}_1 and a smaller part to \mathbf{x}_2 , while ALE attributes the heterogeneity more equally between the two split features. This discrepancy might be explained by the correlation between \mathbf{x}_2 and \mathbf{x}_3 , which particularly affects the two methods based on marginal distributions (that is, PD and SD). Furthermore, Table 5 shows that we tend to obtain shallower trees when we use PINT upfront, while we retain the level of heterogeneity reduction by obtaining similar R_j^2 values for the two interacting features \mathbf{x}_1 and \mathbf{x}_2 .

Thus, PINT reduces the number of features to consider in the interacting feature subset for GADGET depending on the regarded feature effect method. This leads to better results in GADGET in the sense that we only split w.r.t. truly interacting features defined by PINT and receive shallower (and thus more interpretable) trees.

7. Empirical Validation of the PINT Procedure

We now empirically validate PINT to detect interactions.

Hypothesis. It can be shown empirically that the PINT procedure is faithful w.r.t. the type I error and the power of the test as already validated theoretically in Section 5.

10. While the true functional relationship is here defined without noise to purely measure the effect of the spurious interaction, we analyze the effect of adding noise in Appendix F.2




Figure 7: Boxplots showing the distribution of p-values of each feature for different sample sizes and effect methods over all repetitions when PINT is applied (left) and the distribution of feature-wise H-Statistic values for both sample sizes (right).




Figure 8: Boxplots showing the interaction-related heterogeneity reduction I_z per split feature over 30 repetitions when PD, ALE, or SD is used within GADGET. The rows show the results for the two different sample sizes, and the columns indicate if GADGET is used based on all features without using PINT upfront (left) or GADGET is used based on the feature subset resulting from PINT (right).

Experimental Setting. Let $X_1, X_2, X_3, X_4 \sim \mathcal{U}(-1, 1)$ be four independently distributed features, from which we draw 500 observations. The true functional relationship is given by $\mathbf{y} = \beta \mathbf{x}_1 \mathbf{x}_2 + \epsilon$, $\epsilon \sim \mathcal{N}(0, 1)$ with $\beta \in \{0, 0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2, 2.25, 2.5, 2.75, 3\}$. Hence, for $\beta = 0$, no feature has an influence on \mathbf{y} , which is then defined only by the noise term: $\mathbf{y} = \epsilon$. For each setting, we train an SVM based on an RBF kernel with

n	Method	PINT	R_j^2		Number of Subspaces		
			R_1^2 mean(sd)	R_2^2 mean(sd)	Min	Max	Median
300	ALE	no	0.83 (0.05)	0.83 (0.05)	4	10	7.50
300	ALE	yes	0.83 (0.07)	0.82 (0.05)	1	10	7.00
300	PD	no	0.94 (0.02)	0.84 (0.06)	2	7	4.00
300	PD	yes	0.92 (0.03)	0.80 (0.08)	2	5	3.00
300	SD	no	0.93 (0.04)	0.90 (0.05)	2	11	5.00
300	SD	yes	0.93 (0.04)	0.90 (0.05)	2	11	4.50
500	ALE	no	0.83 (0.04)	0.83 (0.06)	4	10	7.00
500	ALE	yes	0.86 (0.04)	0.82 (0.05)	1	11	6.50
500	PD	no	0.94 (0.03)	0.84 (0.06)	2	7	4.00
500	PD	yes	0.93 (0.03)	0.82 (0.08)	2	5	4.00
500	SD	no	0.93 (0.04)	0.90 (0.05)	2	11	5.00
500	SD	yes	0.93 (0.04)	0.90 (0.05)	2	9	5.00

Table 5: Interaction-related heterogeneity reduction per feature for \mathbf{x}_1 and \mathbf{x}_2 by mean (standard deviation) of R_j^2 and minimum, maximum and median number of final subspaces after applying GADGET based on different sample sizes, effect methods and with and without using PINT upfront.

cost parameter $C = 1$ and choose the inverse kernel width based on the data. We apply PINT for all three feature effect methods and with $s = 100$ and $\alpha = 0.05$. We perform the experiments 1000 times to calculate the rejection rate for each feature and setting.

Results. The rejection rate of all four features and all effect methods for $\beta = 0$ align closely with the significance level $\alpha = 0.05$ (see Figure 9). While the power of the test approaches 1 for increasing effect sizes of the interaction between \mathbf{x}_1 and \mathbf{x}_2 for all three methods, the type I error represented by the rejection rates of \mathbf{x}_3 and \mathbf{x}_4 goes to 0. This trend can be explained as follows: The higher β , the higher the influence of \mathbf{x}_1 and \mathbf{x}_2 and thus the more the model focuses on these two features and ignores the two non-influential features. This can also be seen on the p-value distributions, which, for non-influential features such as \mathbf{x}_4 , becomes more skewed towards 1 the higher the interaction effect, as illustrated in Figure 10. Thus, in the most extreme case, the model disregards the non-influential features entirely, resulting in their local effects’ heterogeneity being exactly 0. However, when permuting \mathbf{y} and refitting the model, random effects may be learned for \mathbf{x}_3 and \mathbf{x}_4 . Hence, the calculated values under the null distribution might be higher than the actual risk value, causing the observed artifact in Figure 9 where the type I error is 0 instead of showing a false rejection rate of 5%.

More diverse and complex settings that confirm these observations and empirically validate their correctness are discussed in Appendix F.2.

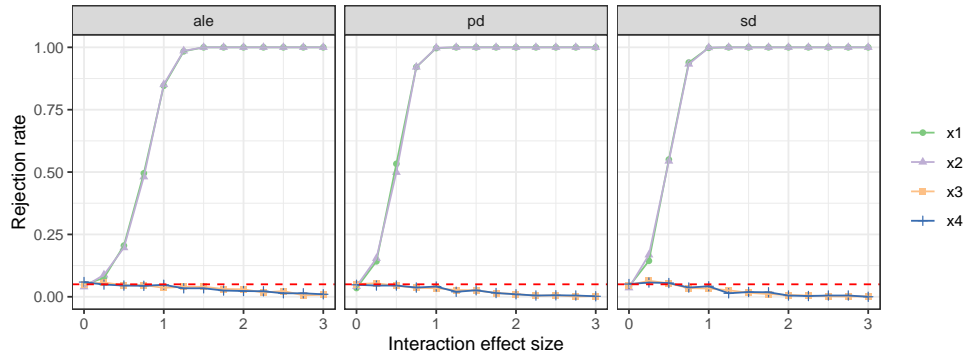


Figure 9: Rejection rates of PINT for 1000 repetitions for PD, ALE, and SD and for varying interaction effect sizes β . For $\beta = 0$, rejection rates correspond to type I errors. For $\beta > 0$, rejection rates of \mathbf{x}_1 and \mathbf{x}_2 correspond to the power of the test, those of \mathbf{x}_3 and \mathbf{x}_4 to type I errors. The red dashed line represents the significance level of $\alpha = 0.05$. Where not visible, the green line (\mathbf{x}_1) is below the purple line (\mathbf{x}_2).

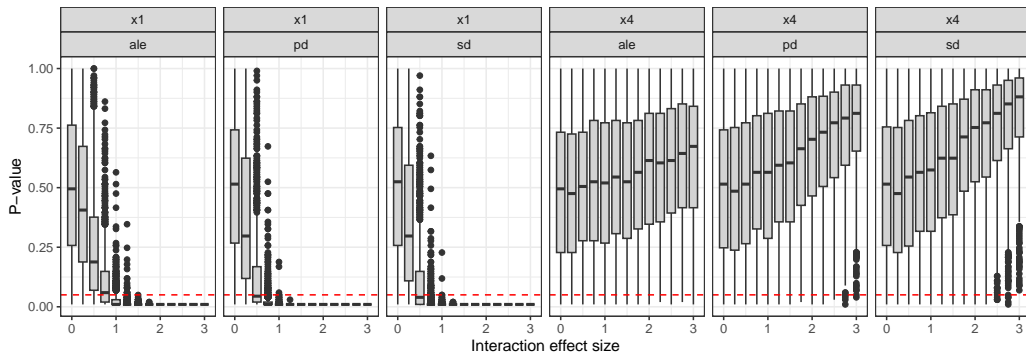


Figure 10: P-value distributions of PINT for 1000 repetitions for \mathbf{x}_1 and \mathbf{x}_4 for all three feature effect methods and for varying interaction effect sizes β . The red dashed line represents the significance level of $\alpha = 0.05$.

8. Real-World Applications

We now show the usefulness of our introduced methodology on real-world applications, revealing insights into the learned regional effects, interactions, and potential biases in the data or model. We present GADGET’s applicability in higher-dimensional settings and offer guidelines for handling such scenarios in Section 9.

8.1 COMPAS Data Set

Due to potential subjective judgement and the resulting bias in the decision making process of the criminal justice system (Blair et al., 2004), ML models have been used to predict

recidivism of defendants to provide a more objective guidance for judges. However, if the underlying training data are biased (e.g., different socioeconomic groups have been treated differently for the same crime in the past), the ML model might learn the underlying data bias, and due to its black-box nature, explanations for its decision-making process and a potential recourse are harder to achieve (Fisher et al., 2019).

We want to use GADGET here to obtain more information on how different characteristics of the defendant and their criminal record influence the risk of recidivism within different subgroups and if “inadmissible” characteristics such as ethnicity or gender cause a different risk evaluation. For our analysis, we use the COMPAS data set to predict the risk of violent recidivism collected by ProPublica (Larson et al., 2016). As Fisher et al. (2019), we choose the three admissible features (1) age of the defendant, (2) number of prior crimes, and (3) if the crime is considered a felony versus misdemeanor, and the two “inadmissible” features (1) ethnicity and (2) gender of the defendant. We use the subset of African-American and Caucasian defendants and apply the data preprocessing steps suggested by ProPublica and applied in Fisher et al. (2019), leaving us with 3373 defendants of the original pool of 4743 defendants. We consider a binary target variable that indicates a high (= 1) or low (= 0) recidivism risk, based on a categorization of ProPublica. We perform our analysis on the entire data set, using a tuned SVM.¹¹

Since the features do not show high correlations, we use ICE and PD for our analysis.¹² Figure 11 shows that the average predicted risk visibly decreases with age, while the predicted risk first increases steeply with the number of previous crimes up to 10 and then decreases slightly for higher values. The PD values for the type of crime do not differ substantially. When considering the two “inadmissible” features, there is on average a slightly higher risk of recidivism for African-American versus Caucasian and female versus male defendants. For all features, we can observe highly differing local effects. Particularly, the heterogeneous ICE curves for age and number of prior crimes indicate feature interactions.

We first apply PINT with PD to define our subset S for GADGET. Using $s = 200$ and $\alpha = 0.05$, we approximate the null distribution using the procedure described in Section 5. All five features are significant w.r.t. the chosen significance level, so we choose $S = Z = \{age, crime, ethnicity, gender, priors_count\}$ for GADGET. We apply GADGET with a maximum depth of 3 and $\gamma = 0.15$. The effect plots for the four resulting regions are shown in Figure 12. GADGET performed the first split according to age and the splits on the second depth according to the number of prior crimes. The total reduction in interaction-related heterogeneity is $R_{Tot}^2 = 0.86$. The highest reduction in heterogeneity is given by age and the number of prior crimes, which interact the most (see Figure 12). For defendants around 20 years of age, the predicted risk tends to be high. However, for defendants with a small number of prior crimes, the predicted risk decreases very quickly with increasing age, reaching a low risk and remaining so for people older than 32 years. In contrast, for defendants with more than four prior crimes, the predicted risk decreases only slowly with increasing age. The regional feature effects of the number of prior crimes show that the interaction-related heterogeneity is small for the subgroup of younger people, while some heterogeneity still remains for defendants older than 32 years of age. Furthermore, the

11. We performed proper model comparison as well to ensure that this is a well-fitting model, at least in comparison to other standard choices.

12. We obtained similar results by using SD instead of PD within GADGET, see Appendix G.




Figure 11: ICE and PD curves of considered features of the COMPAS application example.

interaction-related heterogeneity of the three binary features (ethnicity, gender, and severity of crime) was reduced, indicating an interaction between each of them and age as well as the number of prior crimes. Although the effect on the predicted risk only differs slightly between the categories of the three binary features for older defendants with a lower number of prior crimes and for younger defendants with a high number of prior crimes, greater differences were observed for the other two subgroups. The overall difference in predicted risk for the two inadmissible features of ethnicity and gender seems to be especially high for people over 32 years with a higher number of prior crimes, as well as for people younger than 32 with a lower number of prior crimes. These differences in the predicted risk between subgroups of defendants might be due to a potentially learned bias regarding the ethnicity and gender of the defendant and thus potentially resulting in a more severe predicted risk of recidivism for some defendants than for others.

Note that we applied GADGET to an ML model that is fitted on the COMPAS scores and not directly on COMPAS. Consequently, we are unable to draw conclusions about the learned effects of the underlying commercial black-box model. However, GADGET is model-agnostic and can be applied to any accessible black-box model.

8.2 Bikesharing Data Set

We use the bikesharing data set (James et al., 2022), to predict the hourly counts of rented bikes from the Capital bikeshare system for the years 2011 and 2012. We fit an RF with 10 seasonal and weather features: the day and hour the bike was rented, if the current day is a working day, the season, number of casual bikers, normalized temperature, perceived temperature, wind speed, humidity, and weather situation (categorical, e.g., “clear”).

Again, we first apply PINT with $s = 200$ and $\alpha = 0.05$ on all features to define the interacting feature subset S for GADGET. Since some features—such as season, temperature, and perceived temperature—are correlated, we use ALE for our analysis. Although the features “hour” and “working day” are highly significant, the p-values of all other features are close to 1, indicating that only the heterogeneity of the local effects for hour and workingday is caused by interactions. Hence, we set $S = Z = \{hr, workingday\}$ and apply GADGET with a maximum depth of 3 and $\gamma = 0.15$. GADGET splits once with the binary feature “workingday”, reducing the interaction-related heterogeneity of the two features by




Figure 12: Regional PD plots for features of the COMPAS application after applying GADGET. Grey areas of numeric features and error bars of categorical features indicate the 95% interaction-related heterogeneity interval (see Appendix C.4).

$R_{Tot}^2 = 0.88$. The middle plots of Figure 13 show the regional ALE plots after applying GADGET. During working days and rush hours, prominent peaks are observed, while there is a decline during the noon and afternoon hours. However, on non-working days, the trend is the opposite. This interaction is not visible in the global ALE plot of the feature “hour” (left plot). The interaction-related heterogeneity of the feature “hour” is reduced compared to the global plot, although there is still some apparent variation for working days. From a domain perspective, we might also consider an interaction of the temperature with the hour and working day (as done in Hiabu et al., 2023). Thus, we include the temperature in the feature subsets S and Z and apply GADGET again with the same settings as described above. The feature “workingday” governs the first split. However, for the region of working days, GADGET splits again according to temperature, as shown in the right plot of Figure 13. Although the interaction-related heterogeneity of the feature “temperature” was barely reduced within GADGET ($R_j^2 = 0.03$)—which supports the results of PINT—using temperature in the subset of splitting features Z further reduced the interaction-related heterogeneity of hour by 15%. Thus, feature interactions can be asymmetric, and extending Z based on domain knowledge might be a valid option in some cases.

9. High-Dimensional Real-World Application with Filtering Procedure

We have considered only low-dimensional settings with three to ten features so far. Real-world data sets are often of a higher dimension, posing a challenge for many interpretation methods in general either due to high runtime or potentially overwhelming, incomprehensible outputs (Molnar et al., 2022). Here, we illustrate how this problem can be mitigated for GADGET via a two-step filtering process.

If we can quickly reduce the original feature set to the feature subset which is indeed used by the model and which is interacting, GADGET’s computational complexity then only scales with the number of learned interactions, which can be much smaller than considering all potential interactions of the original feature space. We illustrate this argumentation by the spam data set which contains 57 numeric features to classify 4601 e-mail samples as

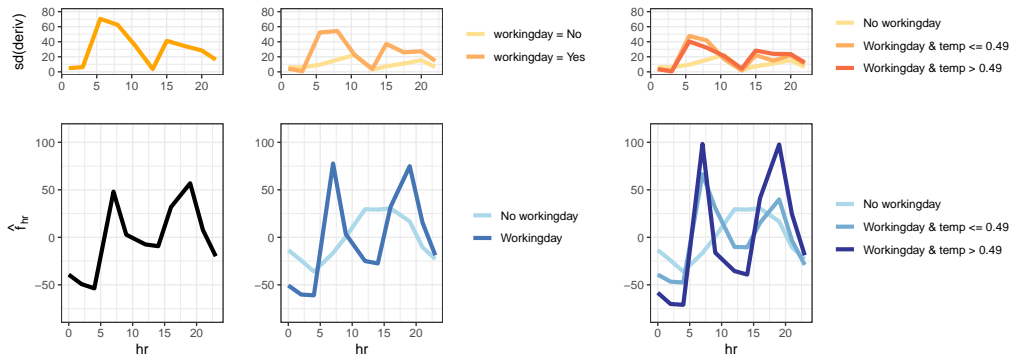


Figure 13: Global (left) and regional ALE plots after applying GADGET for feature hour using $S = Z = \{hr, workingday\}$ (middle) and using $S = Z = \{hr, workingday, temp\}$ (right) of the bikesharing data. The upper plots show the interaction-related heterogeneity as defined in Appendix C.4.

spam or non-spam (Hopkins et al., 1999), based on simple text characteristics. We train an RF with 500 trees, achieving a misclassification error of 4.8% measured by a 5-fold cross-validation, indicating good performance (Hopkins et al., 1999). We apply GADGET-PD since there are only very few correlations present in the feature space.

We suggest a two-step pre-filtering procedure for higher dimensional settings: First, filtering features based on their overall heterogeneity of local effects via Eq. (7). Second, filtering of the remaining features after step 1 for interacting features with PINT according to Algorithm 2. For the first filtering step, we sort the features according to the heterogeneity of local effects measured by Eq. (7) with h being the mean-centered ICE values and \mathcal{A}_g being the entire feature space \mathcal{X} . If the heterogeneity is very small, the shapes of the curves are very similar. Hence, the feature does not interact with other features in the data set and thus does not influence the prediction at all or influences the prediction only via a main effect. Therefore, we can disregard this feature for GADGET.¹³

To decide on the cutoff point of the risk values to group features into one of the two groups: heterogeneous or homogeneous local effects, we visualize the risk values calculated by Eq. (7) in decreasing order by an elbow graph as illustrated in Figure 14. The cutoff point is chosen rather conservatively as the last substantial drop in risk before all remaining values are close to 0.¹⁴ Figure 15 visualizes the respective mean-centered ICE curves for the two features closest to the cutoff point and compares them to the mean-centered ICE curves of the features with the highest and smallest risk values according to Figure 14. This first filtering step leaves 30 features in the group of heterogeneous local effects.

We then apply PINT as a second filtering step, with PD, $s = 200$ permutations, and $\alpha = 0.05$ as described in Section 5. 12 out of the 30 features are detected as significant.

13. In the case of asymmetric feature interactions, we might consider such features in Z but not in S , however, this might require domain knowledge (see the bikesharing example in Section 8).

14. This cutoff point can also be chosen in a data-driven way, e.g., by the percentage of explained variance similarly to principal component analysis.




Figure 14: Risk values according to Eq. (7). Red line marks selected cutoff.




Figure 15: Mean-centered ICE curves for the feature with the highest (left) and lowest (right) risk values and the two features closest to the cutoff point (middle).

We apply GADGET-PD with a relative improvement of $\gamma = 0.25$, a maximum tree depth of 5 and minimum number of observations per node of 100 as early stopping criteria.¹⁵ We define $S = Z$ as the 12 obtained features. We obtain 5 final regions when GADGET-PD is applied. With $R_{Tot}^2 = 0.67$, we removed 67% of the interaction-related heterogeneity within the final regions of all 12 features in S . The visualizations and interpretations of the regional effect plots are provided in Appendix H.

Hence, by using the two filtering steps described above, we reduced the number of plots to examine from 57 to 12. This increases comprehensibility and aids the user to focus on the relevant features and learned relationships.

10. Conclusion and Discussion

We introduced GADGET, a new framework that partitions the feature space into interpretable and distinct regions such that feature interactions of any subset of features are minimized. GADGET can be used with any feature effect method that satisfies the *local decomposability* axiom (Axiom 1). We demonstrated its use with well-known global feature

¹⁵ See Appendix D for general recommendations on specifying these hyperparameters.

effect methods—namely PD, ALE, and SD—and provided visualizations and estimates for regional feature effects and interaction-related heterogeneity. Furthermore, we introduced different measures to quantify and analyze feature interactions through the reduction of interaction-related heterogeneity within GADGET. To define the interacting feature subset, we introduced the novel PINT procedure to detect global feature interactions for any feature effect method used within GADGET.

Our experiments showed that PINT detects the true subset of interacting features and that preselection leads to more meaningful and interpretable results in GADGET. Additionally, feature effect methods based on conditional distributions, such as ALE, tend to produce more stable results compared to those based on marginal distributions, especially when features are highly correlated. Furthermore, due to a different weighting scheme in the decomposition of Shapely values compared to the other considered feature effect methods, higher-order terms are less likely to be detected by GADGET, especially if Shapley values are not recalculated after each partitioning step. This approach might be computationally expensive, which can be seen as a possible limitation. However, recent research has focused on fast approximation techniques of Shapley values (e.g., Lundberg and Lee, 2017; Covert and Lee, 2021; Jethani et al., 2021; Lundberg et al., 2020; Chau et al., 2022; Kolpaczki et al., 2024) and may offer solutions to overcome this limitation.

In general, our proposed method works well if learned feature interactions are not overly local and if observations can be grouped based on feature interactions such that local feature effects within the groups are homogeneous, while different groups exhibit heterogeneous feature effects. This approach helps to avoid the aggregation bias of global feature effect methods, provides deeper insights into learned regional effects, and detects potential biases within subgroups (as illustrated in Section 8). One of the real-world examples also showed that the frequently made assumption of symmetric feature interactions (e.g., in Shapley values) is not always the case (see also Masoomi et al., 2022, for research on asymmetrical feature interactions). Thus, including domain knowledge to define the interacting feature subset Z might sometimes be meaningful.

Finally, GADGET is—as all other interpretation methods—affected by the Rashomon effect due to differently learned feature effects for different fitted models. Thus, the results of GADGET when applied to different models might not be the same, but might provide more insights when compared to each other. Analyzing the influence of the Rashomon effect is out-of-scope in this paper but will be considered in future work. More details and related work to this topic are provided in Appendix I.

Acknowledgments

This work has been partially supported by 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 and by the German Research Foundation (DFG) under the grants 437611051 and 459360854. The authors of this work take full responsibility for its content.

Appendix A. Further Details on Functional ANOVA Decomposition

As stated in Section 2.2, the functional ANOVA provides a decomposition of the prediction function into the main and higher-order effect of all involved features (Sobol', 1993; Stone, 1994; Hooker, 2004, 2007; Li and Rabitz, 2012):

$$\hat{f}(\mathbf{x}) = g_0 + \sum_{j=1}^p g_j(\mathbf{x}_j) + \sum_{j \neq k} g_{jk}(\mathbf{x}_j, \mathbf{x}_k) + \dots + g_{12\dots p}(\mathbf{x}) = g_0 + \sum_{k=1}^p \sum_{\substack{W \subseteq \{1, \dots, p\}, \\ |W|=k}} g_W(\mathbf{x}_W),$$

where g_0 is a constant, $g_j(\mathbf{x}_j)$ denotes the main effect of the j -th feature, $g_{jk}(\mathbf{x}_j, \mathbf{x}_k)$ is the pure two-way interaction effect between features \mathbf{x}_j and \mathbf{x}_k , and so forth. The last component $g_{12\dots p}(\mathbf{x})$ always allows for an exact decomposition. We discuss here the standard and the generalized functional ANOVA decomposition which provide a unique decomposition of the above-stated equation based on different assumptions.

Standard Functional ANOVA Decomposition. The standard functional ANOVA decomposition (Sobol', 1993; Hooker, 2004) provides a unique solution in case of independent features by imposing the vanishing (or strong annihilating) condition (Li and Rabitz, 2012; Rahman, 2014): $\int g_W(\mathbf{x}_W) d\mathbb{P}_{X_j}(X_j) = 0$, $\forall W \neq \emptyset$ and $\forall j \in W$. The vanishing condition implies a) that the component functions $g_W(\mathbf{x}_W)$ “integrate to zero w.r.t. the marginal density of each random variable” in W ; b) the *zero means* property, i.e., $E[g_W(X_W)] = 0$; c) the *orthogonality* property, i.e., $E[g_W(X_W)g_V(X_V)] = 0$ with $\emptyset \neq W \subseteq \{1, \dots, p\}$, $\emptyset \neq V \subseteq \{1, \dots, p\}$ and $W \neq V$; d) that each (lower-dimensional) component function of g_W itself is the zero function. The component functions can be determined recursively by

$$g_W(\mathbf{x}_W) = \int_{\mathbf{x}_{-W}} \left(\hat{f}(\mathbf{x}) - \sum_{V \subset W} g_V(\mathbf{x}_V) \right) d\mathbb{P}_{X_{-W}}(\mathbf{x}_{-W}).$$

Generalized Functional ANOVA Decomposition. In the case of dependent features, the *vanishing condition* does not work, as it does not imply all a)-d) above. Therefore, Hooker (2007) introduced the generalized functional ANOVA decomposition with a *relaxed vanishing* (or weak annihilating) condition: $\int g_W(\mathbf{x}_W) p_X(\mathbf{x}) d\mathbf{x}_j d\mathbf{x}_{-W} = 0$, $\forall j \in W \neq \emptyset$ where p_X is the joint density of X . If this condition holds, then the hierarchical orthogonality condition $\int g_W(\mathbf{x}_W) h_V(\mathbf{x}_V) d\mathbb{P}_X(X) = 0$, $\forall V \subset W$ is satisfied. This again leads to a unique decomposition.

With $p_X(\mathbf{x})$ now being a general probability density with its support being grid-closed and with \hat{f} being square-integrable, the component functions can then be uniquely determined by optimizing Eq. (14).

$$g_W(\mathbf{x}_W) = \arg \min_{h_W \in \mathbb{L}^2(\mathbb{R}^W), W \subseteq \{1, \dots, p\}} \int \left(\sum_{W \subseteq \{1, \dots, p\}} h_W(\mathbf{x}_W) - \hat{f}(\mathbf{x}) \right)^2 d\mathbb{P}_X(\mathbf{x}) \quad (14)$$

The estimation procedure can become computationally expensive with increasing dimensionality. However, being able to decompose a prediction function into its components and hence, being able to interpret main and higher-order effects between features separately

is often desired. Thus, models that are able to directly model these structures have been introduced and extended. Generalized additive models (GAMs), for example, additively decompose the prediction function into the (non-linear) main effects of each feature. Lou et al. (2013) extended GAMs by adding relevant two-way feature interactions to the main effect model. Watanabe et al. (2021) further extend this approach by also allowing feature interactions of a higher order as well as monotonic constraints.

Recent research has focused on computing the decomposition more efficiently, e.g., by developing models that integrate the respective conditions directly in the optimization process through constraints (e.g., Sun et al., 2022) or on purifying the resulting decomposition for specific models within the modeling process (e.g., Lengerich et al., 2020; Hu et al., 2022). However, most approaches require estimating the underlying data distribution which remains challenging (Lengerich et al., 2020), with initial attempts employing adaptive kernel methods (Sun et al., 2022). Notably, finding an appropriate decomposition is only complex in the presence of feature interactions, otherwise, the prediction function can be uniquely decomposed into the main effects using models like the generalized additive model (GAM).

Appendix B. Theoretical Evidence of GADGET

We provide the proofs for the theorems of Sections 4.1 to 4.5. This includes the theoretical foundation of GADGET as well as the proofs of the applicability of the feature effect methods PD, ALE, and SD within the GADGET algorithm by defining respective local feature effect functions that fulfill Axiom 1. We also show in Appendix B.4 that the REPID method which was introduced in Section 3 is a special case of the GADGET algorithm.

B.1 Proof of Theorem 2

Proof Sketch If the function $\hat{f}(\mathbf{x})$ can be decomposed as in Eq. (1) and if Axiom 1 holds for the local feature effect function h , then the main effect of feature \mathbf{x}_j at x_j is cancelled out within the loss function in Eq. (6). Thus, the loss function measures the interaction-related heterogeneity of feature \mathbf{x}_j at x_j , since the variability of local effects in the subspace \mathcal{A}_g are only based on feature interactions between the j -th feature and features in $-j$.

Proof

$$\begin{aligned}
 \mathcal{L}_j(\mathcal{A}_g, x_j) &= \sum_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g} \left(h(x_j, \mathbf{x}_{-j}^{(i)}) - \mathbb{E}[h(x_j, X_{-j}) | \mathcal{A}_g] \right)^2 \\
 &= \sum_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g} \left(g_j(x_j) + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}) - \mathbb{E} \left[g_j(x_j) + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup j}(x_j, X_W) \mid \mathcal{A}_g \right] \right)^2 \\
 &= \sum_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g} \left(\sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}) - \mathbb{E}[g_{W \cup j}(x_j, X_W) | \mathcal{A}_g] \right)^2
 \end{aligned}$$

■

B.2 Proof of Theorem 3

Proof Sketch We show that the theoretical minimum of the objective is $\mathcal{I}(t^*, z^*) = 0$ if no feature in S interacts with any feature in $-Z$. We apply the functional decomposition within the risk function $\mathcal{R}_j(\mathcal{A}_b^{t,z}, \tilde{\mathbf{x}}_j)$. Since we assume that no feature in S interacts with any feature in $-Z$, the second and third term—which contain interactions between these two feature subsets—are zero. Hence, the risk function is only defined by feature interactions between features in S and Z , which are minimized by the objective in Algorithm 1.

Proof If the feature subset Z contains all features interacting with features in S , and hence no feature in $-Z$ interacts with any feature in S , then the risk function for feature \mathbf{x}_j within a subspace $\mathcal{A}_b^{t,z}$ reduces to the variance of feature interactions between feature \mathbf{x}_j and features in Z :

$$\mathcal{R}_j(\mathcal{A}_b^{t,z}, \tilde{\mathbf{x}}_j) = \sum_{\substack{k:k \in \{1, \dots, m\} \\ \wedge x_j^{(k)} \in \mathcal{A}_b^{t,z}}} \sum_{i:\mathbf{x}^{(i)} \in \mathcal{A}_b^{t,z}} \left(\sum_{l=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=l}} g_{W \cup j}(x_j, \mathbf{x}_W^{(i)}) - \mathbb{E}[g_{W \cup j}(x_j, X_W) | \mathcal{A}_b^{t,z}] \right)^2$$

We can further expand the equation above as follows:

$$\begin{aligned} \mathcal{R}_j(\mathcal{A}_b^{t,z}, \tilde{\mathbf{x}}_j) &= \sum_{\substack{k:k \in \{1, \dots, m\} \\ \wedge x_j^{(k)} \in \mathcal{A}_b^{t,z}}} \sum_{i:\mathbf{x}^{(i)} \in \mathcal{A}_b^{t,z}} \left(\left(\sum_{l=1}^{|Z \setminus j|} \sum_{\substack{Z_l \subseteq Z \setminus j, \\ |Z_l|=l}} g_{Z_l \cup j}(x_j, \mathbf{x}_{Z_l}^{(i)}) - \mathbb{E}[g_{Z_l \cup j}(x_j, X_{Z_l}) | \mathcal{A}_b^{t,z}] \right) \right. \\ &+ \left. \left(\sum_{l=2}^{p-1} \sum_{\substack{W \subseteq -j \\ \wedge \exists Z_l \subseteq Z \setminus j: Z_l \subset W \\ \wedge \exists -Z_l \subseteq -Z \setminus j: -Z_l \subset W, \\ |W|=l}} \underbrace{g_{W \cup j}(x_j, \mathbf{x}_{Z_l}^{(i)}, \mathbf{x}_{-Z_l}^{(i)}) - \mathbb{E}[g_{W \cup j}(x_j, X_{Z_l}, X_{-Z_l}) | \mathcal{A}_b^{t,z}]}_{=0} \right) \right)^2 \\ &+ \left(\sum_{l=1}^{|-Z \setminus j|} \sum_{\substack{-Z_l \subseteq -Z \setminus j, \\ |-Z_l|=l}} \underbrace{g_{-Z_l \cup j}(x_j, \mathbf{x}_{-Z_l}^{(i)}) - \mathbb{E}[g_{-Z_l \cup j}(x_j, X_{-Z_l}) | \mathcal{A}_b^{t,z}]}_{=0} \right)^2 \\ &= \sum_{\substack{k:k \in \{1, \dots, m\} \\ \wedge x_j^{(k)} \in \mathcal{A}_b^{t,z}}} \sum_{i:\mathbf{x}^{(i)} \in \mathcal{A}_b^{t,z}} \left(\sum_{l=1}^{|Z \setminus j|} \sum_{\substack{Z_l \subseteq Z \setminus j, \\ |Z_l|=l}} g_{Z_l \cup j}(x_j, \mathbf{x}_{Z_l}^{(i)}) - \mathbb{E}[g_{Z_l \cup j}(x_j, X_{Z_l}) | \mathcal{A}_b^{t,z}] \right)^2 \end{aligned}$$

Since the objective is defined such that it minimizes these interactions for all $j \in S$ by splitting the feature space w.r.t. features in Z , we can split deep enough to achieve $g_{Z_l \cup j}(x_j, \mathbf{x}_{Z_l}^{(i)}) = \mathbb{E}[g_{Z_l \cup j}(x_j, X_{Z_l}) | \mathcal{A}_b^{t,z}]$ for all terms within the sums of the risk function and for all $j \in S$. In other words, the individual interaction effect is equal to the expected interaction effect within a subspace. Thus, the theoretical minimum is $\mathcal{I}(t^*, z^*) = 0$. \blacksquare

B.3 Proof of Theorem 4

Here, we define h to meet Axiom 1 from Section 4.1 for PD where ICE curves represent local feature effects. The i -th ICE curve of feature \mathbf{x}_j can be decomposed as follows:

$$\begin{aligned}
 \hat{f}(\mathbf{x}_j, \mathbf{x}_{-j}^{(i)}) &= \underbrace{g_0}_{\text{constant term}} + \underbrace{g_j(\mathbf{x}_j)}_{\text{main effect of } \mathbf{x}_j} + \underbrace{\sum_{k \in -j} g_k(\mathbf{x}_k^{(i)})}_{\substack{\text{main effect of all other} \\ \text{features in } -j \text{ for observation } i}} \\
 &+ \underbrace{\sum_{k=1}^{p-1} \sum_{W \subseteq -j, |W|=k} g_{W \cup \{j\}}(\mathbf{x}_j, \mathbf{x}_W^{(i)})}_{\substack{(k+1)\text{-order interaction between} \\ \mathbf{x}_j \text{ and features in } -j \text{ for observation } i}} + \underbrace{\sum_{k=2}^{p-1} \sum_{W \subseteq -j, |W|=k} g_W(\mathbf{x}_W^{(i)})}_{\substack{k\text{-order interaction between} \\ \text{features in } -j \text{ for observation } i}}
 \end{aligned}$$

However, this decomposition of the local feature effect of \mathbf{x}_j contains not only feature effects that depend on \mathbf{x}_j , but also other effects (e.g., i -th main effects of features in $-j$), and thus Axiom 1 is not fulfilled by ICE curves.

Proof Sketch By mean-centering ICE curves, constant and feature effects independent of \mathbf{x}_j are canceled out, and thus $\hat{f}^c(\mathbf{x}_j, \mathbf{x}_{-j}^{(i)})$ can be decomposed into the mean-centered main effect of \mathbf{x}_j and the i -th mean-centered interaction effect between \mathbf{x}_j and features in $-j$.

Proof

$$\begin{aligned}
 \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)}) &= \hat{f}(x_j, \mathbf{x}_{-j}^{(i)}) - \mathbb{E} \left[\hat{f}(X_j, \mathbf{x}_{-j}^{(i)}) \right] \\
 &= g_0 + g_j(x_j) + \sum_{k \in -j} g_k(\mathbf{x}_k^{(i)}) + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup \{j\}}(x_j, \mathbf{x}_W^{(i)}) + \sum_{k=2}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_W(\mathbf{x}_W^{(i)}) \\
 &- g_0 - E[g_j(X_j)] - \sum_{k \in -j} g_k(\mathbf{x}_k^{(i)}) - E \left[\sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup \{j\}}(X_j, \mathbf{x}_W^{(i)}) \right] - \sum_{k=2}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_W(\mathbf{x}_W^{(i)}) \\
 &= g_j(x_j) - E[g_j(X_j)] + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup \{j\}}(x_j, \mathbf{x}_W^{(i)}) - E \left[\sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup \{j\}}(X_j, \mathbf{x}_W^{(i)}) \right] \\
 &= \underbrace{g_j^c(x_j)}_{\substack{\text{mean-centered} \\ \text{main effect of } \mathbf{x}_j}} + \underbrace{\sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} g_{W \cup \{j\}}^c(x_j, \mathbf{x}_W^{(i)})}_{\substack{\text{mean-centered interaction effect of} \\ \mathbf{x}_j \text{ with } \mathbf{x}_{-j}^{(i)}}}
 \end{aligned}$$

Thus, Axiom 1 is satisfied by mean-centered ICE curves and can be used as local feature effect h within GADGET. ■

Following from that, the mean-centered PD for feature \mathbf{x}_j at x_j can be decomposed by:

$$f_j^{PD,c}(x_j) = \mathbb{E}[\hat{f}^c(x_j, X_{-j})] = g_j^c(x_j) + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} E \left[g_{W \cup \{j\}}^c(x_j, X_W) \right],$$

which is the mean-centered main effect of feature \mathbf{x}_j and the expected mean-centered interaction effect with feature \mathbf{x}_j at feature value x_j .

B.4 REPID as Special Case of GADGET

The objective function $\mathcal{I}(t, z)$ in Algorithm 1 for $h = \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)})$ is defined by the above loss function $\mathcal{L}_j^{PD}(\mathcal{A}_g, x_j)$ as follows:

$$\mathcal{I}(t, z) = \sum_{j \in S} \sum_{g \in \{l, r\}} \sum_{k: k \in \{1, \dots, m\} \wedge \mathbf{x}_j^{(k)} \in \mathcal{A}_g} \mathcal{L}_j^{PD}(\mathcal{A}_g, \mathbf{x}_j^{(k)})$$

For the special case where we consider one feature of interest that we want to visualize ($S = j$) and all other features as possible split features ($Z = -j$), the objective function of GADGET reduces to:

$$\mathcal{I}(t, z) = \sum_{g \in \{l, r\}} \sum_{k=1}^m \mathcal{L}_j^{PD}(\mathcal{A}_g, \mathbf{x}_j^{(k)}),$$

which is the same objective used within REPID. Thus, for the special case where we choose mean-centered ICE curves as local feature effect method and $S = j$ and $Z = -j$, GADGET is equivalent to REPID.

B.5 Proof of Theorem 5

Here, we show the fulfillment of Axiom 1 from Section 4.1 for ALE.

Proof Sketch The local feature effect method used in ALE is the partial derivative of the prediction function at $\mathbf{x}_j = x_j$. Thus, we define the local feature effect function h by $h(x_j, \mathbf{x}_{-j}^{(i)}) := \frac{\partial \hat{f}(x_j, \mathbf{x}_{-j}^{(i)})}{\partial x_j}$. We can decompose h such that it only depends on main and interaction effects of and with feature \mathbf{x}_j .

Proof

$$\begin{aligned} \frac{\partial \hat{f}(x_j, \mathbf{x}_{-j}^{(i)})}{\partial x_j} &= \frac{\partial \left(g_0 + \sum_{j=1}^p g_j(x_j) + \sum_{j \neq k} g_{jk}(x_j, \mathbf{x}_k^{(i)}) + \dots + g_{12\dots p}(\mathbf{x}^{(i)}) \right)}{\partial x_j} \\ &= \frac{\partial g_j(x_j)}{\partial x_j} + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} \frac{\partial g_{W \cup \{j\}}(\mathbf{x}_j, \mathbf{x}_W^{(i)})}{\partial x_j} \end{aligned}$$

■

Taking the conditional expectation over the local feature effects (partial derivatives) at x_j yields the (conditional) expected (i.e., global) feature effect at x_j :

$$\mathbb{E} \left[\frac{\partial \hat{f}(X_j, X_{-j})}{\partial x_j} \Big| X_j = x_j \right] = \frac{\partial g_j(x_j)}{\partial x_j} + \sum_{k=1}^{p-1} \sum_{\substack{W \subseteq -j, \\ |W|=k}} \mathbb{E} \left[\frac{\partial g_{W \cup j}(X_j, X_W)}{\partial x_j} \Big| X_j = x_j \right]$$

B.6 Proof of Theorem 6

Here, we show the fulfillment of Axiom 1 defined in Section 4.1 for Shapley values, which are the underlying local feature effect in SD plots.

Proof Sketch The local feature effect function in the SD plot is the Shapley value. We define $h(x_j, \mathbf{x}_{-j}^{(i)}) := \phi_j^{(i)}(x_j)$ to be the Shapley value for the i -th local feature effect at a fixed value x_j , which is typically the i -th feature value of \mathbf{x}_j (i.e., $x_j = x_j^{(i)}$). In Eq. (19), we defined $\phi_j^{(i)}(x_j)$ according to Herren and Hahn (2022) which only depends on main and interaction effects of \mathbf{x}_j .

Proof

$$\begin{aligned} \phi_j^{(i)}(x_j) &= \sum_{k=0}^{p-1} \frac{1}{k+1} \sum_{\substack{W \subseteq -j: \\ |W|=k}} \left(\mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] - \sum_{V \subset \{W \cup j\}} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}] \right) \\ &= g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}), \end{aligned}$$

with $g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}) = \mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] - \sum_{V \subset \{W \cup j\}} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}]$.

Hence, we can decompose h such that it only depends on main effects of and interaction effects with feature \mathbf{x}_j . ■

Taking the expectation over the local feature effects $h = \phi_j$ at x_j yields the expected (i.e., global) feature effect of Shapley values at $\mathbf{x}_j = x_j$.

$$\mathbb{E}_{X_W}[\phi_j(x_j)] = g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} \mathbb{E}[g_{W \cup j}^c(x_j, X_W)]$$

Appendix C. Further Characteristics of Feature Effect Methods

In this section, we cover further characteristics of the different feature effect methods used within GADGET. As illustrated for PD in Section 4.3, we also show here for ALE and SD

that the joint feature effect (and possibly the prediction function) within the final regions of GADGET can be approximated by the sum of univariate feature effects. Hence, the joint feature effect can be additively decomposed into the features' main effects within the final regions. Furthermore, we provide an overview on estimates and visualization techniques for the regional feature effects and interaction-related heterogeneity for the different feature effect methods. We also explain how categorical features are handled within GADGET, depending on the underlying feature effect method.

C.1 Decomposability of PD

According to Friedman (2001) the usage of (multivariate) PD functions as additive components in a functional decomposition can recover the prediction function up to a constant. Thus, the prediction function can be decomposed into an intercept g_0 and the sum of mean-centered PD functions similar to the recursive procedure of the standard functional ANOVA decomposition:

$$\hat{f}(\mathbf{x}) = g_0 + \sum_{k=1}^p \sum_{\substack{W \subseteq \{1, \dots, p\}, \\ |W|=k}} \left(f_W^{PD,c}(\mathbf{x}_W) - \sum_{V \subset W} f_V^{PD,c}(\mathbf{x}_V) \right), \quad (15)$$

with $f_W^{PD,c}(\mathbf{x}_W) = f_W^{PD}(\mathbf{x}_W) - \int f_W^{PD}(\mathbf{x}_W) d\mathbf{x}_W$. Tan et al. (2023) note that if the prediction function can be written as a sum of main effects, it can be exactly decomposed by an intercept plus the sum of all mean-centered 1-dimensional PD functions.

Based on this decomposition and if Z contains all features interacting with features in S and if GADGET is applied such that the theoretical minimum of the objective function is reached, then according to Theorem 3, the joint mean-centered PD function $f_{S|\mathcal{A}_g}^{PD,c}$ within each final subspace \mathcal{A}_g can be decomposed into the respective 1-dimensional mean-centered PD functions:

$$f_{S|\mathcal{A}_g}^{PD,c}(\mathbf{x}_S) = \sum_{j \in S} f_{j|\mathcal{A}_g}^{PD,c}(\mathbf{x}_j). \quad (16)$$

Eq. (16) is justified by the assumption that no more interactions between features in S and other features are present in the final regions (Friedman and Popescu, 2008).

Furthermore, if the subset $-S$ is the subset of features that do not interact with any other features (local feature effects are homogeneous), then according to Eq. (16) and Eq. (15), the prediction function $\hat{f}_{\mathcal{A}_g}$ within each final subspace \mathcal{A}_g can be decomposed into the 1-dimensional mean-centered PD functions of all p features plus a constant value g_0 :

$$\hat{f}_{\mathcal{A}_g}(\mathbf{x}) = g_0 + \sum_{j=1}^p f_{j|\mathcal{A}_g}^{PD,c}(\mathbf{x}_j).$$

Thus, depending on how we choose the subsets S and Z and the extent to which we are able to minimize the interaction-related heterogeneity of feature effects by recursively applying Algorithm 1, we might be able to approximate the prediction function by an additive function of main effects of all features within the final regions.

This can also be shown for the simulation example illustrated in Figure 3, where $\hat{f}_{2|\mathcal{A}_g}^{PD,c}(\mathbf{x}_2) = 0$ (the regional effect of feature \mathbf{x}_2 after the split is still 0 and a has low

interaction-related heterogeneity). Instead, the regional effects of \mathbf{x}_1 and \mathbf{x}_3 vary compared to the root node and strongly reduce the interaction-related heterogeneity. Since the regional effects of \mathbf{x}_1 and \mathbf{x}_3 are approximately linear, we can estimate the prediction function within each subspace by

$$\hat{f}_{A_l}(\mathbf{x}) = g_0 + \frac{-3.04}{1.05}\mathbf{x}_1 + \frac{0.99}{0.94}\mathbf{x}_3 = g_0 - 2.9\mathbf{x}_1 + 1.05\mathbf{x}_3$$

and

$$\hat{f}_{A_r}(\mathbf{x}) = g_0 + \frac{3.08}{1.05}\mathbf{x}_1 + \frac{0.89}{0.83}\mathbf{x}_3 = g_0 + 2.93\mathbf{x}_1 + 1.07\mathbf{x}_3,$$

which is a close approximation to the true functional relationship of the data-generating process and thus provides a better understanding of how the features of interest influence the prediction function compared to only considering the global PD plots.

Note that besides the decomposability property in Eq. (16), each global PD of features in S is a weighted additive combination of the final regional PDs. Thus, each global PD can be additively decomposed into regional PD.¹⁶

C.2 Decomposability of ALE

Figure 16 shows the effect plots when GADGET-ALE is applied to the uncorrelated simulation example of Section 4.3 with $S = Z = \{1, 2, 3\}$. The grey curves before the split illustrate the global ALE curves which typically do not show the interaction-related heterogeneity of the underlying local effects. Hence, we added a plot that visualizes this heterogeneity by providing the standard deviation of the underlying derivatives within each interval as a (yellow) curve along the range of \mathbf{x}_j , similar to the idea of Goldstein et al. (2015) for derivative ICE curves. This shows that the local effects for the feature \mathbf{x}_2 are very homogeneous across its entire range, while \mathbf{x}_1 exhibits consistently heterogeneous local effects, and \mathbf{x}_3 displays a high heterogeneity specifically around $\mathbf{x}_3 = 0$. GADGET chooses $\mathbf{x}_3 = -0.003$ as the best split point, significantly reducing the interaction-related heterogeneity among the three features. Thus, the ALE curves in the subspaces more accurately represent the underlying individuals.

Similar to PD plots, ALE plots include an additive recovery and can therefore be decomposed additively into main and interaction effects, as defined in Eq. (17). Apley and Zhu (2020) defined the W -th order ALE function by the pure W -th order (interaction) effect similar to the functional ANOVA decomposition from Eq.(1). Furthermore, Apley and Zhu (2020) show that the ALE decomposition has an orthogonality-like property, which guarantees (similar to the generalized functional ANOVA) that the following decomposition is unique (for details on estimating higher-order ALE functions, see Apley and Zhu (2020)):

$$\hat{f}(\mathbf{x}) = g_0 + \sum_{\substack{W \subseteq \{1, \dots, p\}, \\ |W|=k}} \hat{f}_W^{ALE,c}(\mathbf{x}_W). \quad (17)$$

Furthermore, if Z includes all features interacting with those in S and GADGET achieves the theoretical minimum of the objective function, then according to Theorem 3, the joint

16. This decomposition is not in general true for mean-centered PD curves, since the mean-centering constant might be changed in the partitioning process.

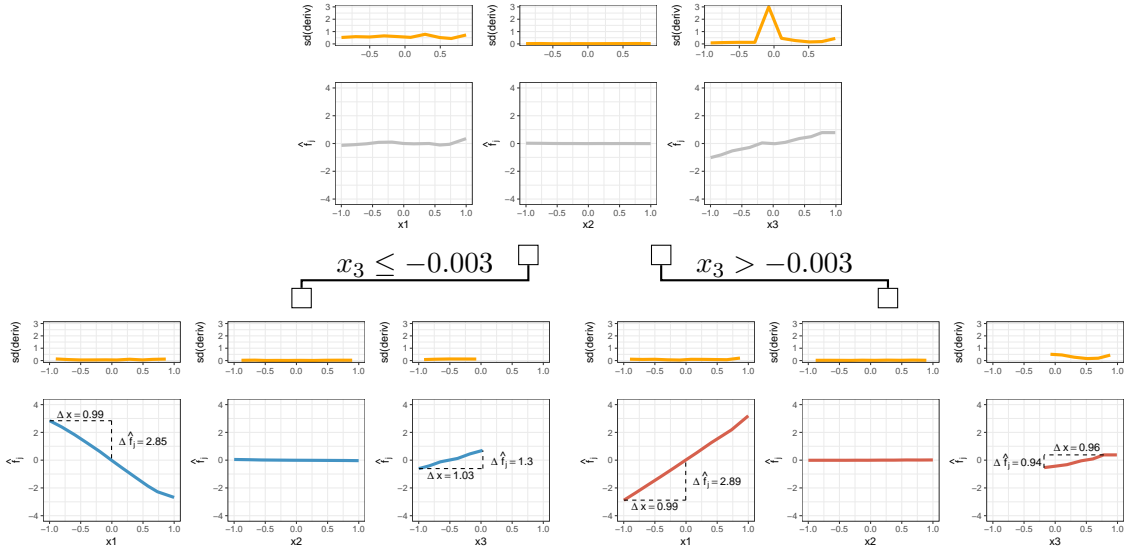


Figure 16: Visualization of applying GADGET with $S = Z = \{1, 2, 3\}$ to derivatives of ALE for the uncorrelated simulation example of Section 4.3 with $Y = 3X_1\mathbb{1}_{X_3>0} - 3X_1\mathbb{1}_{X_3\leq 0} + X_3 + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.09)$. The upper plots show the standard deviation of the derivatives (yellow) and the ALE curves (grey) on the entire feature space, while the lower plots represent the respective standard deviation of the derivatives and regional ALE curves after partitioning the feature space w.r.t. $\mathbf{x}_3 = -0.003$.

mean-centered ALE function $f_{S|\mathcal{A}_g}^{ALE,c}$ within each final subspace \mathcal{A}_g can be decomposed into the 1-dimensional mean-centered ALE functions of features in S . Consequently, with no further interactions between features in S and other features in the final regions, $f_{S|\mathcal{A}_g}^{ALE,c}$ can be uniquely decomposed into the mean-centered main effects of features in S —just as in PD functions (Apley and Zhu, 2020):

$$f_{S|\mathcal{A}_g}^{ALE,c}(\mathbf{x}_S) = \sum_{j \in S} f_{j|\mathcal{A}_g}^{ALE,c}(\mathbf{x}_j) \quad (18)$$

Moreover, let $-S$ be the subset of features that do not interact with any other features. Then, according to Eq. (18) and Eq. (15), the prediction function $\hat{f}_{\mathcal{A}_g}$ within the region \mathcal{A}_g can be decomposed into the 1-dimensional mean-centered ALE functions of all p features, plus some constant value g_0 :

$$\hat{f}_{\mathcal{A}_g}(\mathbf{x}) = g_0 + \sum_{j=1}^p f_{j|\mathcal{A}_g}^{ALE,c}(\mathbf{x}_j).$$

We can again illustrate this decomposition in Figure 16, where \mathbf{x}_2 shows an effect of 0 with low heterogeneity before and after the split. The feature effects of \mathbf{x}_1 and \mathbf{x}_3 show high

heterogeneity before the split, which is almost completely minimized after the split w.r.t. \mathbf{x}_3 . Hence, the resulting regional (linear) ALE curves are representative estimates for the underlying local effects. Therefore, we can approximate the prediction function within each subspace by

$$\hat{f}_{A_l}(\mathbf{x}) = g_0 + \frac{-2.85}{0.99}\mathbf{x}_1 + \frac{1.3}{1.03}\mathbf{x}_3 = g_0 - 2.89\mathbf{x}_1 + 1.26\mathbf{x}_3$$

and

$$\hat{f}_{A_r}(\mathbf{x}) = g_0 + \frac{2.89}{0.99}\mathbf{x}_1 + \frac{0.94}{0.96}\mathbf{x}_3 = g_0 + 2.92\mathbf{x}_1 + 0.98\mathbf{x}_3.$$

Particularities of ALE Estimation. As seen for the continuous feature \mathbf{x}_3 in the simulation example presented here, abrupt interactions (“jumps”)¹⁷ might be difficult to estimate for models that learn smooth effects, such as NNs (used here) or SVMs—especially when compared to models such as decision trees. Hence, depending on the model, these type of feature interactions can lead to very high partial derivatives in a region around the “jump” point instead of a high partial derivative at exactly the one specific “jump” point (here: $\mathbf{x}_3 = 0$), thus leading to non-reducible heterogeneity (see upper right plot in Figure 16). The standard deviation of the derivatives of \mathbf{x}_3 are very high in the region around and not exactly at $\mathbf{x}_3 = 0$. This interaction-related heterogeneity should be (almost) completely reduced when splitting w.r.t. $\mathbf{x}_3 = 0$. However, high values may persist if the model did not perfectly capture this kind of interaction. To account for this issue in the estimation and partitioning process within GADGET, we use the following procedure for continuous features: In the two new subspaces after a split, if the derivatives of feature values close to the split point vary at least twice as much (measured by the standard deviation) as those of other observations within each subspace, we replace these derivatives close to the split point with values drawn from a normally distributed random variable where mean and variance are estimated by the derivatives of the remaining observation within each subspace.

C.3 Decomposability of SD

Recalculation Versus No Recalculation of Shapley Values. In Section 4.5, we argued that Shapley values must be recalculated after each partitioning step to ensure that each new subspace to receive SD effects in the final subspaces that are representative of the underlying main effects within each subspace. Meanwhile, the unconditional expected value (i.e., the feature interactions on the entire feature space) are minimized without recalculating the conditional expected values.

The difference in the final feature effects within the subspaces is illustrated when comparing the left lower plot of Figure 4 (split without recalculation) with the respective plots of feature \mathbf{x}_1 of Figure 17. Without recalculation, the effect of feature \mathbf{x}_1 is still regarded as an interaction effect between \mathbf{x}_1 and \mathbf{x}_3 , and hence only half of the joint interaction effect is assigned to \mathbf{x}_1 (i.e., the respective slope within the regions is 1.5 and -1.5 instead of 3 and -3), and the other half of the joint interaction effect is assigned to \mathbf{x}_3 . When Shapley values are recalculated after the first partitioning step within each subspace, we can see in

17. With abrupt interaction, we mean interactions that lead to an abrupt change of the influence of one feature (\mathbf{x}_1) based on the influence of another feature at a specific (“jump”) point ($\mathbf{x}_3 = 0$) like the feature interaction between \mathbf{x}_1 and \mathbf{x}_3 in the here presented simulation example: $Y = 3X_1\mathbb{1}_{X_3>0} - 3X_1\mathbb{1}_{X_3\leq 0} + X_3 + \epsilon$ with $\epsilon \sim \mathbb{N}(0, 0.09)$.

Figure 17 that no more interactions are present between \mathbf{x}_1 and \mathbf{x}_3 within each subspace, due to the split w.r.t. \mathbf{x}_3 . Hence, the effect of \mathbf{x}_1 is recognized as the main effect, with a slope closely matching the true functional relationship. Furthermore, recalculation also reduces the heterogeneity of feature effects of \mathbf{x}_3 due to interactions with \mathbf{x}_1 .

Note: If the feature we use for partitioning the feature space ($z \in Z$) coincides with the features of interest (S), then the Shapley values should be recalculated in Algorithm 1 to find the best split point (at least, if we choose the approach with recalculation after each partitioning step). The reason is that if $z \in S$, we also want to reduce the interaction-related heterogeneity within z that is not accounted for if we do not recalculate the Shapley values within the new subspace. For example, in Figure 17, we split according to \mathbf{x}_3 , which is also a feature of interest (here: $\{3\} \in S$). If we do not recalculate the Shapley values for \mathbf{x}_3 within the splitting process, then the sum of the risk of any two subspaces for \mathbf{x}_3 will always be approximately the same as the risk of the parent node, and thus the heterogeneity reduction for \mathbf{x}_3 (see regional plots in Figure 17) is not considered in the objective of Algorithm 1.

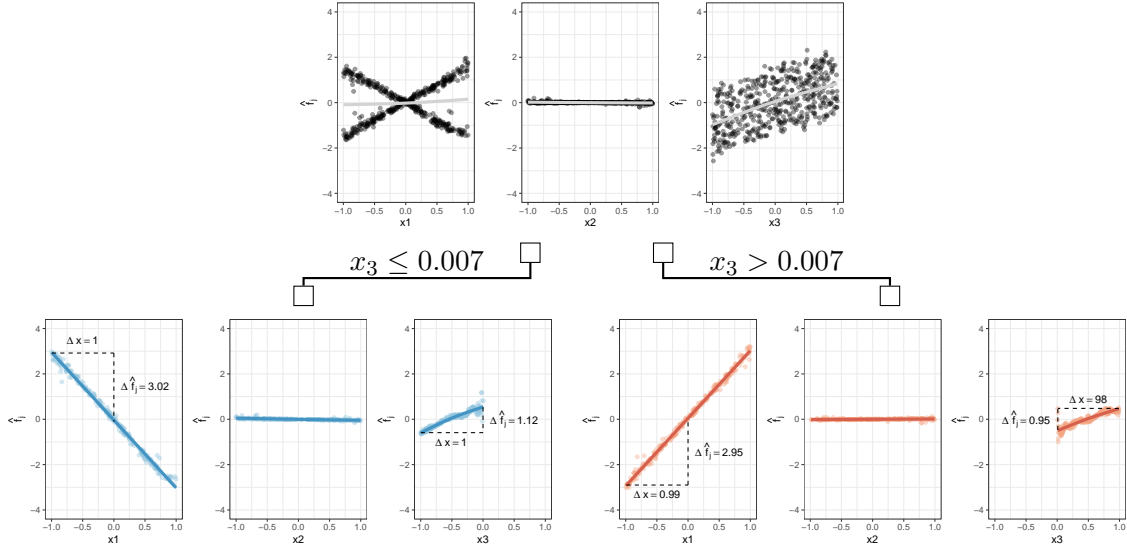


Figure 17: Visualization of applying GADGET with $S = Z = \{1, 2, 3\}$ to Shapley values of the uncorrelated simulation example of Section 4.3 with $Y = 3X_1 \mathbb{1}_{X_3 > 0} - 3X_1 \mathbb{1}_{X_3 \leq 0} + X_3 + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.09)$. The upper plots show the Shapley values and the global estimated SD curve on the entire feature space, while the lower plots represent the respective Shapley values and regional SD curves after partitioning the feature space w.r.t. $\mathbf{x}_3 = 0.007$.

Decomposition. Herren and Hahn (2022) showed that the Shapley value of feature \mathbf{x}_j at feature value x_j can be decomposed according to the functional ANOVA decomposition into

main and interaction effects¹⁸:

$$\begin{aligned}\phi_j^{(i)}(x_j) &= \sum_{k=0}^{p-1} \frac{1}{k+1} \sum_{\substack{W \subseteq -j: \\ |W|=k}} \left(\mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] - \sum_{V \subset \{W \cup j\}} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}] \right) \\ &= g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}).\end{aligned}\quad (19)$$

Thus, in the case of interventional (i.e., marginal-based) Shapley values, Eq. 19 are given by weighted PD functions. Hence, if the global SD feature effect as defined in Eq. (11) is considered, the same decomposition rules as defined for PD plots apply. In other words, if Z contains all features that interact with features in S and if GADGET is applied such that the theoretical minimum of the objective function is reached, then according to Theorem 3, the following decomposition in 1-dimensional global SD effect functions of features in S holds:

$$f_{S|\mathcal{A}_g}^{SD}(\mathbf{x}_S) = \sum_{j \in S} f_{j|\mathcal{A}_g}^{SD}(\mathbf{x}_j). \quad (20)$$

If all features containing heterogeneous effects (feature interactions) are included in the subset S , and the subset Z consists of all features that interact with features in S , then according to Eq. (20) and Eq. (15), the prediction function $\hat{f}_{\mathcal{A}_g}$ within the region \mathcal{A}_g can be uniquely decomposed into the 1-dimensional global SD effect functions of all p features, plus some constant value g_0 :

$$\hat{f}_{\mathcal{A}_g}(\mathbf{x}) = g_0 + \sum_{j=1}^p f_{j|\mathcal{A}_g}^{SD}(\mathbf{x}_j).$$

Again, we can derive this decomposition from Figure 16 in the same way we did for PD and ALE plots. Hence, we can approximate the prediction function within each subspace by $\hat{f}_{\mathcal{A}_l}(\mathbf{x}) = g_0 - 3.02\mathbf{x}_1 + 1.12\mathbf{x}_3$ and $\hat{f}_{\mathcal{A}_r}(\mathbf{x}) = g_0 + 2.98\mathbf{x}_1 + 1.03\mathbf{x}_3$.

Equivalence of SD and Mean-Centered PD. According to Herren and Hahn (2022), the Shapley value $\phi_j^{(i)}(x_j)$ of the i -th observation at $\mathbf{x}_j = x_j$ can be decomposed as in Eq. (19):

$$\begin{aligned}\phi_j^{(i)}(x_j) &= \sum_{k=0}^{p-1} \frac{1}{k+1} \sum_{\substack{W \subseteq -j: \\ |W|=k}} \left(\mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] - \sum_{V \subset \{W \cup j\}} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}] \right) \\ &= g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}),\end{aligned}$$

with $g_{W \cup j}^c(x_j, \mathbf{x}_W^{(i)}) = \mathbb{E}[\hat{f}(x_j, X_{-j}) | X_W = \mathbf{x}_W^{(i)}] - \sum_{V \subset \{W \cup j\}} \mathbb{E}[\hat{f}(X) | X_V = \mathbf{x}_V^{(i)}]$. As in Eq. (11), the global feature effect (SD) of feature \mathbf{x}_j at x_j is then defined by

$$f_j^{SD}(x_j) = \mathbb{E}_{X_W}[\phi_j] = g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} \mathbb{E}[g_{W \cup j}^c(x_j, X_W)]$$

18. Similar decompositions have been introduced in Hiabu et al. (2023) and Bordt and von Luxburg (2023).

Hence, if Theorem 3 is satisfied, if the joint global SD effect of features in S can be decomposed into the univariate SD effects as in Eq. (20), and if the interventional approach for Shapley calculation is used, then all feature interactions are zero, and the global SD effect of feature \mathbf{x}_j at x_j is given by

$$\begin{aligned} f_j^{Shap}(x_j) &= g_j^c(x_j) + \sum_{k=1}^{p-1} \frac{1}{k+1} \sum_{W \subseteq -j: |W|=k} \mathbb{E}[g_{W \cup j}^c(x_j, X_W)] \\ &\stackrel{\text{T. 3}}{=} g_j^c(x_j) \\ &= \mathbb{E}[\hat{f}(x_j, X_{-j})] - \mathbb{E}[\hat{f}(X)], \end{aligned}$$

which is equivalent to the mean-centered PD of feature \mathbf{x}_j at x_j .

C.4 Overview on Estimates and Visualizations

We provide here an overview on the estimates and visualization techniques for PD, ALE, and SD within GADGET that we introduced in Sections 4.3-4.5.

Local Effect. The local effect h for a feature \mathbf{x}_j at feature value x_j used within GADGET is estimated by

- **PD:** mean-centered ICE $\hat{h}^{(i)} = \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)})$
- **ALE:** partial derivatives estimated by prediction differences within k -th interval $\hat{h}^{(i)} = \hat{f}(z_{k-1,j}, \mathbf{x}_{-j}^{(i)}) - \hat{f}(z_{k-1,j}, \mathbf{x}_{-j}^{(i)})$ where $x_j \in]z_{k-1,j}, z_{k,j}]$
- **SD:** Shapley value $\hat{h}^{(i)} = \hat{\phi}_j^{(i)}$

Regional Effect. The feature effect for a feature \mathbf{x}_j at feature value x_j within a subspace/region \mathcal{A}_g of GADGET is estimated by

- **PD:** mean-centered regional PD $\hat{f}_{j|\mathcal{A}_g}^{PD,c}(x_j) = \frac{1}{|N_g|} \sum_{i \in N_g} \hat{f}^c(x_j, \mathbf{x}_{-j}^{(i)})$ with N_g being the index set of all $i : \mathbf{x}^{(i)} \in \mathcal{A}_g$.
- **ALE:** regional ALE $\hat{f}_{j|\mathcal{A}_g}^{ALE}(x_j) = \sum_{k=1}^{k_j(x_j)} \frac{1}{|N_g(k)|} \sum_{i \in N_g(k)} [\hat{f}(z_{k,j}, \mathbf{x}_{-j}^{(i)}) - \hat{f}(z_{k-1,j}, \mathbf{x}_{-j}^{(i)})]$ with $N_g(k)$ being the index set of all $i : x_j \in]z_{k-1,j}, z_{k,j}] \wedge \mathbf{x}^{(i)} \in \mathcal{A}_g$.
- **SD:** regional SD $\hat{f}_{j|\mathcal{A}_g}^{SD}(x_j)$ is estimated by fitting a GAM on $\{x_j^{(i)}, \hat{\phi}_j^{(i)}\}_{i: \mathbf{x}^{(i)} \in \mathcal{A}_g}$.

The regional effect for feature \mathbf{x}_j is visualized for all $x_j \in \mathcal{A}_g$. Therefore, the respective GAM curve is plotted in the case of SD, while we linearly interpolate between the grid-wise/interval-wise estimates of PD/ALE to receive the regional effect curves.

Interaction-Related Heterogeneity. The interaction-related heterogeneity for a feature \mathbf{x}_j at feature value x_j within a subspace/region \mathcal{A}_g of GADGET is estimated by the loss function in Eq. (6), which quantifies the variance of local effects at x_j and is visualized by

- **PD:** 95% interval for (mean-centered) regional PD $[\hat{f}_{j|\mathcal{A}_g}^{PD,c}(x_j) \pm 1.96 \cdot \sqrt{\hat{\mathcal{L}}_j^{PD}(\mathcal{A}_g, x_j)}]$.

- **ALE**: standard deviation of local effects $\sqrt{\hat{\mathcal{L}}_j^{ALE}(\mathcal{A}_g, x_j)}$.
- **SD**: Shapley values are recalculated within each region \mathcal{A}_g and plotted with the fitted GAM for the regional SD effect to visualize the variation of local feature effects aka interaction-related heterogeneity.

For each feature $j \in S$, we generate one figure showing the regional effect curves of all final regions we obtain after applying GADGET. For PD, the regional effect curves are accompanied with intervals showing how much interaction-related heterogeneity remains in the underlying local effects (see, e.g., Figures 12). For ALE, a separate plot visualizes the interaction-related heterogeneity via the standard deviation of local effects, which is inspired by the derivative ICE plots of Goldstein et al. (2015) (see, e.g., Figure 13). For SD, the Shapley values that were recalculated conditioned on each subspace \mathcal{A}_g are visualized with the regional effect curve (see, e.g., Figure 22).

Note that we can also visualize the non-centered PD ($\hat{f}_{j|\mathcal{A}_g}^{PD}$) instead of the mean-centered PD, which might provide more insights regarding interpretation. However, the interaction-related heterogeneity must be estimated by the mean-centered ICE curves to only represent heterogeneity induced by feature interactions (see Appendix B.3).

C.5 Handling of Categorical Features

In this section, we will summarize the particularities of categorical features. Compared to numeric features, we have a limited number of K values (categories). Hence, compared to numeric features, we find split points by dividing the K categories into the two new subspaces. Since GADGET is based on the general concept of a CART (decision tree) algorithm (Breiman et al., 1984) that can handle categorical features, the splitting itself follows the same approach as for a common CART algorithm. If categorical features occur in S , the calculation of the objective function remains unchanged, and the handling of categorical features depends only on the underlying feature effect method:

PD Plot. Compared to numeric features, the grid points for categorical features are limited to the number of categories. Otherwise, the calculation of the loss and the risk (see, Eq. 9 and Eq. 7) works the same as for numeric features.¹⁹

ALE Plot. ALE builds intervals based on quantiles for numeric features to calculate prediction differences between neighboring interval borders for all observations falling within this interval. For binary features, the authors solve this as follows: for all observations falling in each of the categories, the prediction difference when changing it to the other category is calculated. For more categories, they suggest a sorting algorithm.²⁰ Hence, we still receive the needed derivatives for GADGET for each category to calculate the loss and risk function for GADGET.

SD Plot. Compared to numeric features, the x-axis of the SD plot is a grid of size K . For each of these grid points (categories), the Shapley values for the observations belonging to

19. Since we calculate the loss point-wise at each grid point and sum it up over all grid points, the order of the category does not make a difference for the objective of GADGET.

20. For more information, we refer to Apley and Zhu (2020).

the specific category are calculated. Hence, instead of a spline to quantify the expected value in Eq. (11), we use the arithmetic mean within each category (similar to PD plots) and, thus, sum up the variance of Shapley values for each category over all categories (within the respective subspace).

In general, if we apply GADGET and split w.r.t. a categorical feature such that only one category is present within a subspace (e.g., we split the feature sex such that all individuals are male in one resulting subspace), then the interaction-related heterogeneity vanishes to zero since only an additive shift for the feature sex is left in this subspace.

Note: If a categorical feature \mathbf{x}_j is not only considered for splitting ($j \in Z$) but is also a feature of interest ($j \in S$), the different splitting possibilities of categories of \mathbf{x}_j prompt recalculation for ALE, since derivatives are only calculated for pre-sorted neighboring categories. In our implementations, we only split w.r.t. the pre-sorted categories and considered them as integer values to reduce the computational burden of the calculations.

Appendix D. Definition and Specification of Stop Criteria for GADGET

The question of how many partitioning steps should be performed depends generally on the underlying research question. If the user is more interested in reducing the interaction-related heterogeneity as much as possible, they might split rather deeply, depending on the complexity of interactions learned by the model. However, this might lead to many regions that are more challenging to interpret. If the user is more interested in a small number of regions, they might prefer a shallow tree, thus reducing only the heterogeneity of the features that interact the most.

D.1 Definition of Stop Criteria

We propose the following stopping criteria to control the number of partitioning steps in GADGET: First, we can use common decision tree hyperparameters such as the tree depth or the minimum number of observations per leaf node. Alternatively, an early stop mechanism can be applied based on the interaction-related heterogeneity reduction. According to our proposed split-wise measure, we further split only if the relative improvement is at least $\gamma \in [0, 1]$ times the total relative heterogeneity reduction of the previous split: $\gamma \times \frac{\sum_{j \in S} (\mathcal{R}_j(\mathcal{A}_P, \tilde{\mathbf{x}}_j)) - \mathcal{I}(\hat{t}, \hat{z})}{\sum_{j \in S} (\mathcal{R}_j(\mathcal{X}, \tilde{\mathbf{x}}_j))}$. Another option is to stop splitting once a pre-defined total reduction of heterogeneity (R_{Tot}^2) is achieved. Generally, higher γ values and lower thresholds for R_{Tot}^2 result in fewer partitioning steps, and vice versa.

D.2 Recommendations to Specify Stop Criteria

GADGET performs binary splits, making it easier to minimize abrupt interactions (i.e., interactions where categorical or discretized features are involved, see also Section C.2). Conversely, minimizing smooth interactions (i.e., interactions between continuous features) may require multiple splits. We recommend a maximum tree depth between 3 and 6 to maintain some degree of interpretability and setting γ (the minimum relative risk reduction for a split) between 0.1 and 0.25 to further limit the number of terminal regions. For abrupt interactions (with few interactions), the risk reduction for the first split is usually quite high.

Therefore, to find further feature interactions, we recommend to set γ towards the lower end of the recommended range (e.g., 0.1 to 0.2). In the case of smooth interactions, the risk reduction for the first split is usually rather low compared to subgroup-specific feature interactions. Therefore, setting γ to a small value leads to a vanishing improvement which prevents the desired stop mechanism. However, the value should also not be too high to prevent the first split. In summary, a tree depth of 3 to 6 and γ between 0.1 and 0.25 worked well in our examples. To precisely quantify feature interactions or to completely minimize them, more splits and a smaller γ may be needed, possibly using R_{Tot}^2 as a stopping criterion.

Appendix E. Theoretical Evidence of PINT

Here, we provide the proofs to the Theorems defined in Section 5.

E.1 Proof of Theorem 7

Proof Sketch In this proof we show that the type I error $\mathbb{P}(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha})$ of the test is constrained by $\alpha + \epsilon$ with $\epsilon = \max_{t>0} [\mathbb{P}(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) > t) - \mathbb{P}(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > t)]$.

Proof The type I error probability is

$$\begin{aligned} & \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha}\right) \\ &= \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha}\right) + \left[\mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha}\right) - \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha}\right)\right] \\ &\leq \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > z_{1-\alpha}\right) + \max_{t>0} \left[\mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) > t\right) - \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > t\right)\right] \\ &\leq \alpha + \epsilon, \end{aligned}$$

where the last step follows from the definition of $z_{1-\alpha}$ and independence of permutations. ■

E.2 Proof of Theorem 8

Proof Sketch Here, we show that the type II error of the PINT procedure strives towards 0. Recall that a risk of 0 means that no interactions are learned by the model. If Condition (12) holds—i.e., the risk under the permuted data set is less than every $\epsilon > 0$ with high probability—the type II error of the PINT procedure strives towards 0. The last step follows from the assumption that feature interactions have been learned by the model fitted on the original data set. It follows that the power of the test strives towards 1.

Proof Condition (12) implies

$$\mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n^\Pi}, \tilde{\mathbf{x}}_j, \mathcal{X}) > \epsilon\right) \rightarrow 0$$

for every $\epsilon > c > 0$ and thus $\mathbb{P}(z_{1-\alpha} \geq c) \rightarrow 0$. The type II error probability is

$$\begin{aligned}
 & \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) \leq z_{1-\alpha}\right) \\
 & \leq \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) \leq z_{1-\alpha}, z_{1-\alpha} \geq c\right) + \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) \leq z_{1-\alpha}, z_{1-\alpha} < c\right) \\
 & \leq \mathbb{P}\left(z_{1-\alpha} \geq c\right) + \mathbb{P}\left(\mathcal{R}_j(\hat{f}_{\mathcal{D}_n}, \tilde{\mathbf{x}}_j, \mathcal{X}) < c\right) \\
 & \rightarrow 0.
 \end{aligned}$$

■

Appendix F. Further Empirical Validation of PINT

F.1 Example for Failure of PINT

Note that if the condition (12) in Section 5.2 is not fulfilled, the type I error is not constrained by the significance level α . This can be illustrated with the following example: Let $X_1, X_2, X_4 \sim \mathcal{U}(-1, 1)$ and $X_3 = \exp(X_2) + \delta$ with $\delta \sim \mathcal{N}(0, 0.01)$. We draw 500 observations and define the true functional relationship by $y = \beta \mathbf{x}_1 \exp(\mathbf{x}_2) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 1)$ and $\beta \in \{0, 0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2, 2.25, 2.5, 2.75, 3\}$. We again fit an SVM with RBF kernel and apply PINT using the same hyperparameter settings as before. We repeat the experiment 1000 times. In this setting, \mathbf{x}_3 almost perfectly reflects the influence of $\exp(\mathbf{x}_2)$ and therefore it is possible to learn the simpler linear $\beta \mathbf{x}_1 \mathbf{x}_3$ instead of the true more complex non-linear interaction effect. Hence, it is very likely that $\hat{f}(\mathbf{x})$ does not accurately approximate $f(\mathbf{x})$ since $\hat{f}(\mathbf{x})$ most likely learns an interaction between \mathbf{x}_1 and the non-influential feature (in $f(\mathbf{x})$) \mathbf{x}_3 . Hence, this can lead to a higher type I error for \mathbf{x}_3 than α as illustrated in Figure 18.

F.2 Further Empirical Validation of PINT

Here, we provide further evaluations and empirical validation of using PINT to select truly interacting features. The experiments in this section are based on the simulation study of Section 6.3 where we analyzed the robustness of PINT concerning potential spurious interactions and compared it to the H-Statistic. We now analyze in more detail the sensitivity of PINT for the three feature effect methods considering varying correlations of the features, a linear and non-linear true underlying function, differing effect sizes of the interaction effect and how the noise term of the true underlying function influences the rejection rate of PINT. For all experiments, we consider drawing 500 observations of the considered random variables and an SVM based on an RBF kernel with the hyperparameters Section 6.3. More specifically we consider the following combination of different simulation settings:

Linear and non-linear functional relationships. We consider two different true underlying functions. One with linear feature effects (as in Section 6.3) and one more complex one with non-linear effects. Both of them include feature interactions between \mathbf{x}_1 and \mathbf{x}_2 and




Figure 18: Rejection rates of PINT for 1000 repetitions for all three feature effect methods and for varying interaction effect sizes β . For $\beta = 0$, all rejection rates correspond to type I errors. For $\beta > 0$, the rejection rates of \mathbf{x}_1 and \mathbf{x}_2 correspond to the power of the test, those of \mathbf{x}_3 and \mathbf{x}_4 to type I errors. The red dashed line in the left plot represents the significance level of $\alpha = 0.05$.

potential spurious interactions. We argue that the linear setting that we considered so far is rather simple leading to a model fit that is (almost) perfect without too much model uncertainty. Therefore, considering a more complex setting might lead to a worse and less certain model fit.

- Linear true functional relationship: $\mathbf{y} = f_{lin}(\mathbf{x}) + \epsilon_k$ with $f_{lin}(\mathbf{x}) = \mathbf{x}_1 + \mathbf{x}_2 + \mathbf{x}_3 - \beta \mathbf{x}_1 \mathbf{x}_2$ and $k \in \{1, 2\}$.
- Non-linear true functional relationship: $\mathbf{y} = f_{nlin}(\mathbf{x}) + \epsilon_k$ with $f_{nlin}(\mathbf{x}) = \mathbf{x}_1^2 + \mathbf{x}_2^3 + \exp(\mathbf{x}_3) - \beta \mathbf{x}_1 \mathbf{x}_2 + \epsilon_k$ with $k \in \{1, 2\}$.

Correlations and spurious interactions. We consider four features $X_1, X_2, X_4 \sim \mathcal{U}(-1, 1)$ and $X_3 = X_2 + \delta$ with $\delta \sim \mathcal{N}(0, 0.09)$ for the setting defined in Section 6.3 with a high linear correlation of $\rho_{23} \approx 0.9$ between \mathbf{x}_3 and \mathbf{x}_2 . This will very likely lead to a spurious interaction between \mathbf{x}_1 and \mathbf{x}_3 as shown in the analysis of the mentioned section. To investigate how sensitive PINT is concerning this phenomenon, we consider a second smaller correlation which reduces the risk of spurious interactions. Therefore, we take into account a medium correlation by defining X_3 as a combination of X_2 and a uniformly distributed variable (similarly to Section 6.1): $X_3 = 0.37X_2 + 0.63u$ with $u \sim \mathcal{U}(-1, 1)$ which leads to a linear correlation of $\rho_{12} \approx 0.5$. As a plausibility check and baseline, we consider all features to be independent and thus $X_3 \sim \mathcal{U}(-1, 1)$.

Noise term. Regarding the noise term ϵ_k we take into consideration two alternatives: no noise term, i.e., $\epsilon_1 = 0$. This is the setting we used in Section 6.3. Therefore, the noise did not influence the modeling approach and thus did also not influence the interaction detection with PINT. To analyze the influence of a noise term on PINT, we consider one that accounts

for ten percent of the standard deviation of $f(X)$ (similar to the other simulations in this paper), i.e., $\epsilon_2 \sim N(0, 0.01 \cdot \text{var}(f(X)))$.

Interaction effect size. Both functional relationships are also defined for different coefficient values β for the interaction terms. In Section 6.3 we chose $\beta = 2$. The analysis showed that the interaction can be detected very well. Hence, we take into consideration further smaller values for β to see how it affects the power of the test. Therefore, we consider $\beta \in \{0.75, 1, 1.25, 1.5, 1.75, 2\}$.

For each combination of the described settings, we perform 1000 repetitions. The rejection rates over these repetitions for the different settings of the linear functional relationship are shown in Figure 19. It shows that when SD is used within PINT for the uncorrelated setting the type I and type II error are always 0 independent of the interaction effect size. This can be similarly observed for PD, however, for ALE the interaction is not detected for $\beta = 0.75$. Also, the other settings show a higher power for PD and SD. A reason for this might be the definition of the intervals for ALE since the results for ALE are more sensitive concerning the size of the intervals than PD for the number of grid points. It is also observable that the power of the test is in most cases slightly higher for the settings without noise than with noise, however, the differences are rather small. In general, the interaction effect is harder to detect for smaller effect sizes if the correlation increases. The type I error does never exceed the significance level $\alpha = 0.05$, also not for \mathbf{x}_3 which might have been considered in a spurious interaction with \mathbf{x}_1 for the medium and high correlation settings.

For the non-linear setting (see Figure 20), we can observe type I and II errors of 0 for the settings without and with medium correlations if PD or SD is used within PINT for all interaction effect sizes. Compared to the linear setting, a type I error that exceeds the significance level is observable for the high correlation settings. The type I error is therefore particularly high for increasing effect sizes when PD is considered. Learning non-linear effects leads to more oscillating behavior when extrapolating into unseen regions of the feature space. As discussed and illustrated in Section 4.6, this problem is particularly severe for PD plots leading to high heterogeneity and, therefore, could also affect PINT.

Appendix G. GADGET-SD Results for COMPAS

COMPAS Data Set. In addition to the results of GADGET based on PD presented in Section 8, we also applied GADGET with the same settings based on SD.

The effect plots for the four resulting regions are shown in Figures 21 and 22. GADGET based on SD finds the same first split as for PD. The second split is also executed according to the number of prior crimes, but the split value is lower than for PD (at 2.5 instead of 4.5). The total interaction-related heterogeneity reduction ($R_{Tot}^2 = 0.87$) is also similar. Note that Shapley values explain the difference between the actual and average prediction. Hence, SD plots are centered, while Figure 12 shows the uncentered regional PD plots.

Appendix H. Further Details on Higher Dimensional Settings

In Section 9, we consider the spam data set as an example of a high-dimensional data set. The data set contains 4601 samples of e-mails of 57 numeric features and a binary target

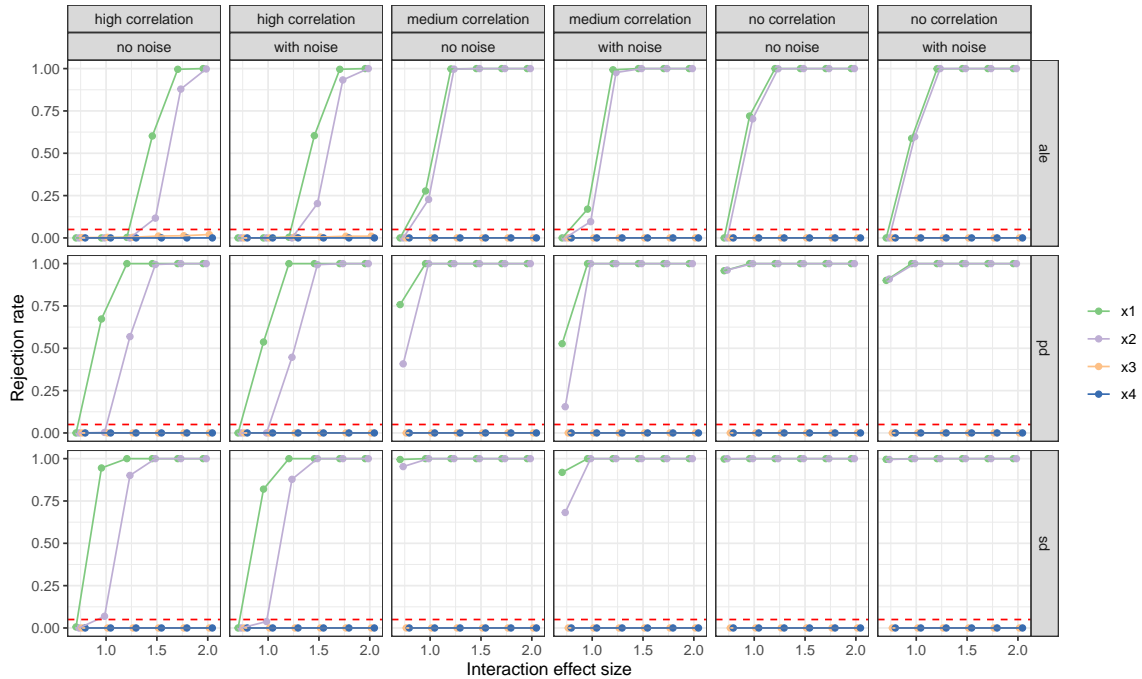


Figure 19: Rejection rates across all runs for the linear functional relationship and for each feature effect method shown in the rows. The columns are defined by the correlation between \mathbf{x}_2 and \mathbf{x}_3 and if a noise term is included or not. For each of these settings varying interaction effect sizes between 0.75 and 2 are considered.

(i.e., spam (1) or non-spam (0)). The features *capitalLong*, *capitalAve*, and *capitalTotal* reflect the length of the longest, the average length and the total length of uninterrupted sequences of capital letters in an e-mail, respectively. The remaining 48 features refer to the percentage of the respective word or character appearing within an e-mail.

While we explained in Section 9 the filtering procedure to handle high-dimensional data, we show here the results after applying the filtering steps and GADGET-PD on the spam data set. Therefore, we apply GADGET on the 12 remaining features after the filtering procedure described in Section 9. Figure 23 shows the regional effects plots for 3 of the 12 features. The first split feature chosen by GADGET is *charDollar* which indicates the percentage of the dollar sign within an e-mail. A high number is an indicator for spam e-mails. This region represented by the yellow curves is highly influenced by the number of occurrences of the word *edu*. The combination of a lower number of dollar signs with small values for *capitalLong*, decreases the probability of being classified as spam to most likely being classified as non-spam. However, in combination with a higher frequency of the words *remove* or *credit* the classification as spam becomes more likely. For a small number of dollar signs but higher values of *capitalLong*, also the frequencies of the word *remove* and of *charExclamation* are used to define the remaining three regions in Figure 23.

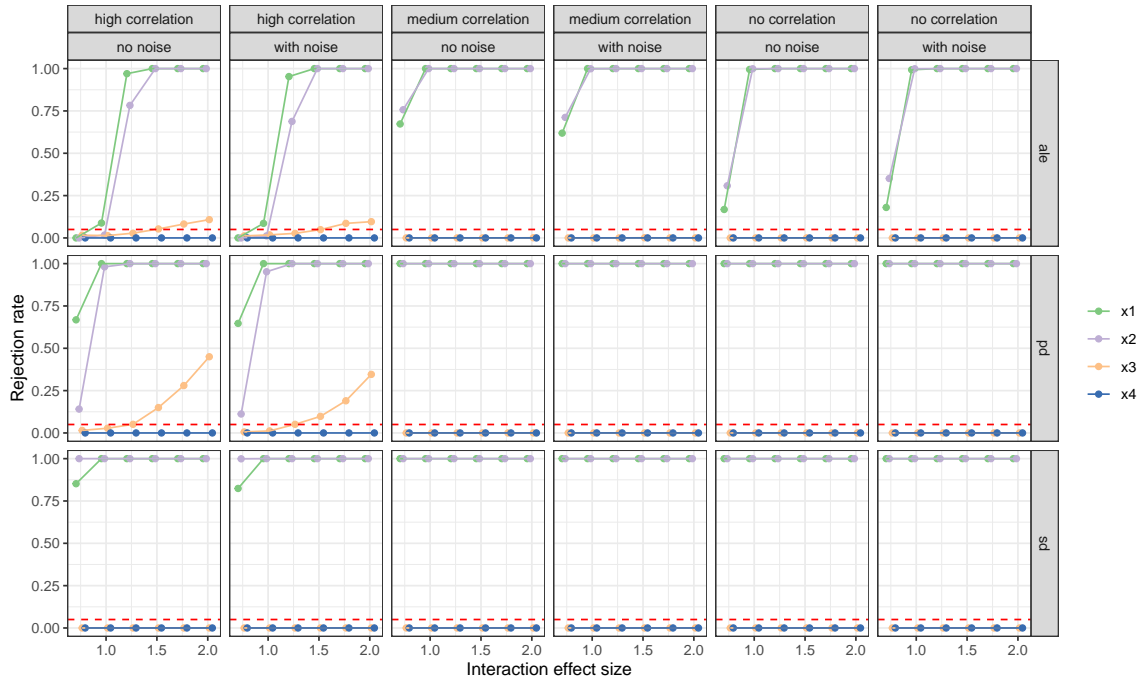


Figure 20: Rejection rates across all runs for the non-linear functional relationship and for each feature effect method shown in the rows. The columns are defined by the correlation between \mathbf{x}_2 and \mathbf{x}_3 and if a noise term is included or not. For each of these settings varying interaction effect sizes between 0.75 and 2 are considered.




Figure 21: Global SD curves and Shapley values of considered features of the COMPAS application example.

Appendix I. Rashomon Effect

The Rashomon effect in ML states that there are often various models of a function class with almost equally good performance that are constructed differently (Breiman, 2001; Semenova et al., 2022; Molnar et al., 2022). This set of models is called the Rashomon set.

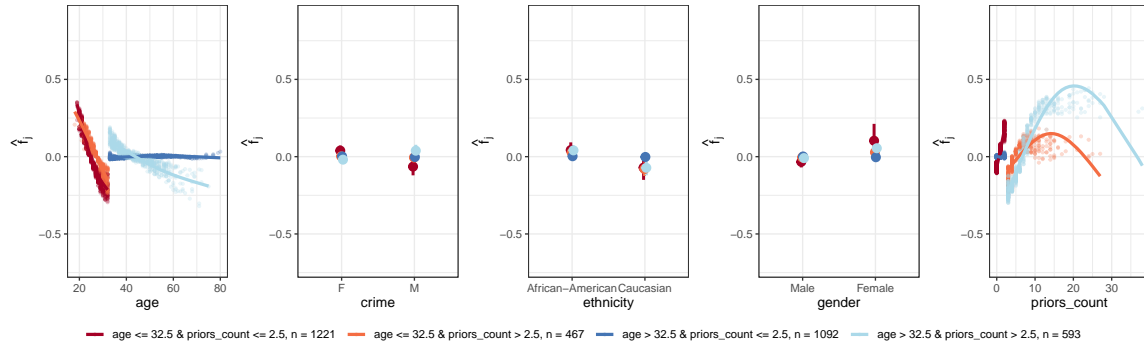


Figure 22: Regional SD plots for considered features of the COMPAS application after applying GADGET. Shapley values within each region are recalculated and visualize the interaction-related heterogeneity within each region.

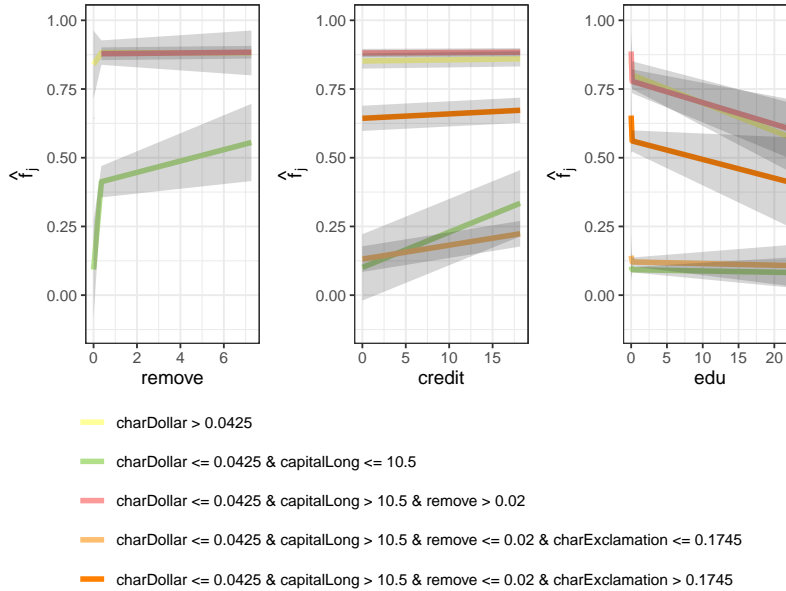


Figure 23: Regional PD plots for *remove*, *credit* and *edu* of the spam data set after applying GADGET when only interacting features are selected for S and Z according to the pre-filtering steps. The grey areas indicate the 95% interaction-related heterogeneity interval as defined in Appendix C.4.

Since these models are constructed differently, they might identify different patterns in the data and rely on different features, which can lead to different results in the interpretation methods. For example, Müller et al. (2023) evaluate empirically, how the Rashomon effect influences different feature attribution methods. Their findings provide empirical support for both the disagreement problem of different interpretation methods and the sensitivity of these methods concerning their hyperparameter configurations. Hence, the Rashomon effect

can also appear on the level of the interpretation method. A common example is seen in counterfactual explanations, where different counterfactuals may provide the same desired change in the outcome but based on different changes in the feature space. To address this issue, Dandl et al. (2020) recommended to consider multiple counterfactuals instead of only one counterfactual. Furthermore, the decomposition of the prediction function of a fitted model \hat{f} into effects of varying orders is typically not unique. As a result, many different exact decompositions are possible and can lead to differing interpretations (Lengerich et al., 2020). One possibility to obtain a unique decomposition is the functional ANOVA decomposition. However, this decomposition might still differ across other models, as it is specifically calculated for one particular fitted model. Admittedly, the Rashomon effect remains an important unsolved problem in interpretable ML and a further discussion is outside of the scope of this work, but we refer the reader to Fisher et al. (2019) or Donnelly et al. (2023) for potential approaches to deal with it.

Appendix J. Marginal- and Conditional-based Feature Effect Methods

Feature effect methods can be grouped into approaches based on marginal or conditional distributions. While the PD plot is exemplary for the former and the M plot as well as the ALE plot are based on the latter, Shapley values can be defined either way.

While marginal-based approaches are usually simple to define and efficiently computable, they are known to extrapolate into unseen regions of the feature space when features are highly correlated, which might lead to misleading interpretations as illustrated in Section 4.6 for PD plots (Hooker, 2007; Molnar et al., 2022). Conditional-based approaches solve the extrapolation problem by considering the true underlying data distribution (Aas et al., 2021). However, they require estimating the conditional data distribution, which is challenging, particularly in high-dimensional settings. The estimation influences the quality of the derived interpretation (Aas et al., 2021; Sundararajan and Najmi, 2020; Janzing et al., 2020). While both approaches come with different merits and weaknesses, the derived interpretation often differs as well. For example, PD plots visualize the influence of a feature of interest on the prediction function irrespective of the underlying feature correlations. In contrast, the M plot accounts for the correlation structure and thus also takes into account the influence of the features correlated with the feature of interest. Therefore, the two approaches answer different questions: Marginal-based approaches primarily offer insights into the inner workings of the ML model without explicitly considering the underlying data correlations, making them particularly useful for model audit and debugging. Conditional-based approaches are used to derive interpretations that reflect the given data distribution and thus are preferable for questions related to model inference (Chen et al., 2020; Freiesleben et al., 2024; Watson, 2022).

To conclude, Chen et al. (2020) argue that the marginal-based approach is not generally wrong and can be useful if we want to derive explanations that are true to the model, while the conditional-based approach should be used when we want to extract interpretations that are true to the data. The framework that we introduce in this paper works with both marginal-based (e.g., PD plots as illustrated in Section 4.3) as well as conditional-based (e.g., ALE as shown in Section 4.4) approaches. Hence, the user needs to decide on a suitable feature effect method depending on their desired interpretation.

References

- Kjersti Aas, Martin Jullum, and Anders Løland. Explaining individual predictions when features are dependent: More accurate approximations to Shapley values. *Artificial Intelligence*, 298:103502, September 2021.
- André Altmann, Laura Toloşi, Oliver Sander, and Thomas Lengauer. Permutation importance: a corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347, 04 2010.
- Daniel W. Apley and Jingyu Zhu. 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, 2020.
- Irene V Blair, Charles M Judd, and Kristine M Chapleau. The influence of afrocentric facial features in criminal sentencing. *Psychological science*, 15(10):674–679, 2004.
- Sebastian Bordt and Ulrike von Luxburg. From Shapley values to generalized additive models and back. In *International Conference on Artificial Intelligence and Statistics*, pages 709–745. PMLR, 2023.
- Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231, 2001.
- Leo Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, 1984.
- Matthew Britton. Vine: Visualizing statistical interactions in black box models. *arXiv preprint arXiv:1904.00561*, 2019.
- Giuseppe Casalicchio, Christoph Molnar, and Bernd Bischl. Visualizing the feature importance for black box models. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part I 18*, pages 655–670. Springer, 2019.
- Siu Lun Chau, Robert Hu, Javier Gonzalez, and Dino Sejdinovic. Rkhs-shap: Shapley values for kernel methods. In *Advances in Neural Information Processing Systems*, volume 35, pages 13050–13063, 2022.
- Hugh Chen, Joseph D. Janizek, Scott Lundberg, and Su-In Lee. True to the model or true to the data? *arXiv preprint arXiv:2006.16234*, 2020.
- Hugh Chen, Ian C. Covert, Scott M. Lundberg, and Su-In Lee. Algorithms to estimate Shapley value feature attributions. *Nature Machine Intelligence*, 5(6):590–601, 2023. Publisher: Nature Publishing Group UK London.
- Ian Covert and Su-In Lee. Improving kernelshap: Practical Shapley value estimation using linear regression. In *International Conference on Artificial Intelligence and Statistics*, pages 3457–3465. PMLR, 2021.

- Susanne Dandl, Christoph Molnar, Martin Binder, and Bernd Bischl. Multi-objective counterfactual explanations. In *International Conference on Parallel Problem Solving from Nature*, pages 448–469. Springer, 2020.
- Jon Donnelly, Srikar Katta, Cynthia Rudin, and Edward P Browne. The rashomon importance distribution: Getting rid of unstable, single model-based variable importance. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Michael D Ernst. Permutation methods: a basis for exact inference. *Statistical Science*, pages 676–685, 2004.
- Aaron Fisher, Cynthia Rudin, and Francesca Dominici. 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, 2019.
- Timo Freiesleben, Gunnar König, Christoph Molnar, and Álvaro Tejero-Cantero. Scientific inference with interpretable machine learning: Analyzing models to learn about real-world phenomena. *Minds and Machines*, 34(3):32, 2024.
- Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001.
- Jerome H. Friedman and Bogdan E. Popescu. Predictive learning via rule ensembles. *The Annals of Applied Statistics*, 2(3):916–954, 2008.
- Alex Goldstein, Adam Kapelner, Justin Bleich, and Emil Pitkin. Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics*, 24(1):44–65, January 2015.
- Andreas Henelius, Kai Puolamäki, Henrik Boström, Lars Asker, and Panagiotis Papapetrou. A peek into the black box: exploring classifiers by randomization. *Data mining and knowledge discovery*, 28:1503–1529, 2014.
- Andreas Henelius, Antti Ukkonen, and Kai Puolamäki. Finding statistically significant attribute interactions. *arXiv preprint arXiv:1612.07597*, 2016.
- Julia Herbinger, Bernd Bischl, and Giuseppe Casalicchio. Repid: Regional effect plots with implicit interaction detection. In *International Conference on Artificial Intelligence and Statistics*, pages 10209–10233. PMLR, 2022.
- Andrew Herren and P. Richard Hahn. Statistical aspects of shap: Functional anova for model interpretation. *arXiv preprint arXiv:2208.09970*, 2022.
- Munir Hiabu, Joseph T Meyer, and Marvin N Wright. Unifying local and global model explanations by functional decomposition of low dimensional structures. In *International Conference on Artificial Intelligence and Statistics*, pages 7040–7060. PMLR, 2023.
- Giles Hooker. Discovering additive structure in black box functions. In *Proceedings of the tenth ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, pages 575–580, 2004.

- Giles Hooker. Generalized functional anova diagnostics for high-dimensional functions of dependent variables. *Journal of Computational and Graphical Statistics*, 16(3):709–732, September 2007.
- Mark Hopkins, Erik Reeber, George Forman, and Jaap Suermondt. Spambase. UCI Machine Learning Repository, 1999.
- Linwei Hu, Jie Chen, Vijayan N Nair, and Agus Sudjianto. Surrogate locally-interpretable models with supervised machine learning algorithms. *arXiv preprint arXiv:2007.14528*, 2020.
- Linwei Hu, Jie Chen, and Vijayan N. Nair. Using model-based trees with boosting to fit low-order functional anova models. *arXiv preprint arXiv:2207.06950*, 2022.
- Gareth James, Daniela Witten, Trevor Hastie, and Rob Tibshirani. *ISLR2: Introduction to Statistical Learning, Second Edition*, 2022. R package version 1.3-2.
- Dominik Janzing, Lenon Minorics, and Patrick Blöbaum. Feature relevance quantification in explainable ai: A causal problem. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 2907–2916. PMLR, 2020.
- Neil Jethani, Mukund Sudarshan, Ian Connick Covert, Su-In Lee, and Rajesh Ranganath. Fastshap: Real-time Shapley value estimation. In *International Conference on Learning Representations*, 2021.
- Patrick Kolpaczki, Viktor Bengs, Maximilian Muschalik, and Eyke Hüllermeier. Approximating the Shapley value without marginal contributions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 13246–13255, 2024.
- Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How we analyzed the compas recidivism algorithm. *ProPublica*, 2016. URL <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.
- Benjamin Lengerich, Sarah Tan, Chun-Hao Chang, Giles Hooker, and Rich Caruana. Purifying interaction effects with the functional anova: An efficient algorithm for recovering identifiable additive models. In *International Conference on Artificial Intelligence and Statistics*, pages 2402–2412. PMLR, 2020.
- Genyuan Li and Herschel Rabitz. General formulation of hdmr component functions with independent and correlated variables. *Journal of Mathematical Chemistry*, 50(1):99–130, January 2012.
- Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.
- Wei-Yin Loh. Regression tress with unbiased variable selection and interaction detection. *Statistica sinica*, pages 361–386, 2002.

- Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. Accurate intelligible models with pairwise interactions. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, KDD '13, pages 623–631, New York, NY, USA, August 2013. Association for Computing Machinery. ISBN 9781450321747. doi: 10.1145/2487575.2487579. URL <https://doi.org/10.1145/2487575.2487579>.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1): 56–67, 2020.
- Aria Masoomi, Davin Hill, Zhonghui Xu, Craig P Hersh, Edwin K. Silverman, Peter J. Castaldi, Stratis Ioannidis, and Jennifer Dy. Explanations of black-box models based on directional feature interactions. In *International Conference on Learning Representations*, 2022.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- Lucas Mentch and Giles Hooker. Formal hypothesis tests for additive structure in random forests. *Journal of Computational and Graphical Statistics*, 26(3):589–597, 2017.
- Christoph Molnar, Gunnar König, Julia Herbinger, Timo Freiesleben, Susanne Dandl, Christian A Scholbeck, Giuseppe Casalicchio, Moritz Grosse-Wentrup, and Bernd Bischl. General pitfalls of model-agnostic interpretation methods for machine learning models. In *xxAI-Beyond Explainable AI: International Workshop, Held in Conjunction with ICML 2020, July 18, 2020, Vienna, Austria, Revised and Extended Papers*, pages 39–68. Springer, 2022.
- Christoph Molnar, Gunnar König, Bernd Bischl, and Giuseppe Casalicchio. Model-agnostic feature importance and effects with dependent features: a conditional subgroup approach. *Data Mining and Knowledge Discovery*, pages 1–39, 2023.
- Sebastian Müller, Vanessa Toborek, Katharina Beckh, Matthias Jakobs, Christian Bauckhage, and Pascal Welke. An empirical evaluation of the Rashomon effect in explainable machine learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 462–478. Springer, 2023.
- Art B. Owen. Variance components and generalized Sobol’ indices. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):19–41, 2013.
- Sharif Rahman. A generalized anova dimensional decomposition for dependent probability measures. *SIAM/ASA Journal on Uncertainty Quantification*, 2:670–697, January 2014.

- Christian A. Scholbeck, Giuseppe Casalicchio, Christoph Molnar, Bernd Bischl, and Christian Heumann. Marginal effects for non-linear prediction functions. *Data Mining and Knowledge Discovery*, Feb 2024.
- Lesia Semenova, Cynthia Rudin, and Ronald Parr. On the existence of simpler machine learning models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1827–1858, 2022.
- Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.
- Margaret A Shipp, Ken N Ross, Pablo Tamayo, Andrew P Weng, Jeffery L Kutok, Ricardo CT Aguiar, Michelle Gaasenbeek, Michael Angelo, Michael Reich, Geraldine S Pinkus, et al. Diffuse large b-cell lymphoma outcome prediction by gene-expression profiling and supervised machine learning. *Nature medicine*, 8(1):68–74, 2002.
- Ilya M Sobol'. Sensitivity estimates for nonlinear mathematical models. *Mathematical Modelling and Computational Experiment*, 1:407–414, 1993.
- Kacper Sokol and Peter Flach. Explainability fact sheets: A framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 56–67, 2020.
- Daria Sorokina, Rich Caruana, Mirek Riedewald, and Daniel Fink. Detecting statistical interactions with additive groves of trees. In *Proceedings of the 25th International Conference on Machine Learning*, pages 1000–1007, 2008.
- Clemens Stachl, Florian Pargent, Sven Hilbert, Gabriella M Harari, Ramona Schoedel, Sumer Vaid, Samuel D Gosling, and Markus Bühner. Personality research and assessment in the era of machine learning. *European Journal of Personality*, 34(5):613–631, 2020.
- Charles J. Stone. The use of polynomial splines and their tensor products in multivariate function estimation. *The Annals of Statistics*, 22(1):118 – 171, 1994.
- Erik Štrumbelj and Igor Kononenko. Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41:647–665, 2014.
- Xingzhi Sun, Ziyu Wang, Rui Ding, Shi Han, and Dongmei Zhang. Puregam: Learning an inherently pure additive model. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1728–1738, 2022.
- Mukund Sundararajan and Amir Najmi. The many Shapley values for model explanation. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9269–9278, 2020.
- Sarah Tan, Giles Hooker, Paul Koch, Albert Gordo, and Rich Caruana. Considerations when learning additive explanations for black-box models. *Machine Learning*, 112(9):3333–3359, 2023.

- Akihisa Watanabe, Michiya Kuramata, Kaito Majima, Haruka Kiyohara, Kondo Kensho, and Kazuhide Nakata. Constrained generalized additive 2 model with consideration of high-order interactions. In *2021 International Conference on Electrical, Computer and Energy Technologies (ICECET)*, pages 1–6, 2021. doi: 10.1109/ICECET52533.2021.9698779.
- David S Watson. Conceptual challenges for interpretable machine learning. *Synthese*, 200(2):65, 2022.
- David S Watson and Marvin N Wright. Testing conditional independence in supervised learning algorithms. *Machine Learning*, 110(8):2107–2129, 2021.
- Xiaohang Zhang, Yuan Wang, and Zhengren Li. Interpreting the black box of supervised learning models: Visualizing the impacts of features on prediction. *Applied Intelligence*, 51(10):7151–7165, October 2021.