
WeightedSHAP: analyzing and improving Shapley based feature attributions

Yongchan Kwon
Columbia University
New York, NY, 10027
yk3012@columbia.edu

James Zou
Stanford University and Amazon AWS
Stanford, CA, 94305
jamesz@stanford.edu

Abstract

Shapley value is a popular approach for measuring the influence of individual features. While Shapley feature attribution is built upon desiderata from game theory, some of its constraints may be less natural in certain machine learning settings, leading to unintuitive model interpretation. In particular, the Shapley value uses the same weight for all marginal contributions—i.e. it gives the same importance when a large number of other features are given versus when a small number of other features are given. This property can be problematic if larger feature sets are more or less informative than smaller feature sets. Our work performs a rigorous analysis of the potential limitations of Shapley feature attribution. We identify simple settings where the Shapley value is mathematically suboptimal by assigning larger attributions for less influential features. Motivated by this observation, we propose WeightedSHAP, which generalizes the Shapley value and learns which marginal contributions to focus directly from data. On several real-world datasets, we demonstrate that the influential features identified by WeightedSHAP are better able to recapitulate the model’s predictions compared to the features identified by the Shapley value.

1 Introduction

Explaining how a feature impacts a model prediction is a crucial question in machine learning (ML) as it provides a deeper understanding of how the model behaves and what insights have been extracted from data. In many real-world applications, it has been increasingly common to deploy complicated models such as a deep neural network model or a random forest to achieve high predictability. However, it often comes with a cost of unintuitive interpretations, and it naturally calls for a principled and practical attribution method. The goal of this work is to quantify the contribution of individual features to a particular prediction, also known as the attribution problem.

Lundberg and Lee [2017] proposed a model-agnostic attribution method, SHapley Additive exPlanations (SHAP), based on the Shapley value from economics. Supported by theoretical properties that the Shapley value satisfies, SHAP has been a popular method in the attribution literature [Janzing et al., 2020, Sundararajan and Najmi, 2020]. For instance, Frye et al. [2020] and Aas et al. [2021] developed the SHAP algorithms for dependent features, and Heskes et al. [2020] and Wang et al. [2021] proposed a rigorous framework in causal inference settings. One of the practical problems to use SHAP is the heavy computational costs, and there have been works on improving computational efficiency [Covert and Lee, 2021, Lundberg et al., 2020, Jethani et al., 2021a]. In addition, SHAP has been deployed to various applied scientific research [Lundberg et al., 2018, Janizek et al., 2021, Qiu et al., 2022].

The Shapley value, a mathematical basis of SHAP, is a simple average of the marginal contributions that quantify the average change in a coalition function when a feature of interest is added from a

subset of features with a given coalition size. There are different version of the marginal contributions by the coalition size, and the Shapley value takes a uniform weight to summarize the influence of a feature. This uniform weight arises due to the efficiency axiom of the Shapley value, which requires the sum of attributions to equal the original model prediction, but it is often problematic because some may be more informative than others. As we will show in Section 3, the Shapley value is not optimal to sort features in order of influence on a model prediction.

Our contributions. While SHAP is widely used for feature attribution, its limitations are still not rigorously understood. We first show the suboptimality of the Shapley value through an analysis of the marginal contributions. We identify a key limitation of the Shapley value in that it assigns uniform weights to marginal contributions in computing the attribution score. We show that this can lead to attribution mistakes when different marginal contributions have different signal and noise. Motivated by this analysis, we propose WeightedSHAP, a generalization of the Shapley value which is more flexible. WeightedSHAP uses a weighted average of marginal contributions where the weights can be learned from the data. On several real-world datasets, our experiments demonstrate that WeightedSHAP is better able to identify influential features that recapitulate a model’s prediction compared to a standard SHAP. WeightedSHAP is a simple modification of SHAP and is easy to adapt from existing SHAP implementations.

1.1 Related works

Model interpretation There are mainly two types of model interpretation depending on the quantity to be accounted for; global and local interpretations. The global interpretation is to explain the impact of a feature on a prediction model across the entire dataset [Lipovetsky and Conklin, 2001, Breiman, 2001, Owen, 2014, Broto et al., 2020, Zhao and Hastie, 2021, Bénard et al., 2022]. For instance, for a decision tree model, Breiman et al. [2017] measures the total decrease of node impurity at node split by a feature of interest as an impact. In contrast, the local interpretation is to explain the impact of a feature on a particular prediction value [Lundberg and Lee, 2017, Chen et al., 2018, Janzing et al., 2020, Lundberg et al., 2020, Heskes et al., 2020, Jethani et al., 2021a]. For a deep neural network model, a gradient-based method uses the gradient evaluated at a particular input sample as an impact [Simonyan et al., 2013, Sundararajan et al., 2017, Ancona et al., 2017, Selvaraju et al., 2017, Adebayo et al., 2018]. Our work studies the local interpretation problem with a focus on a marginal contribution-based method, which we review in Section 2. The marginal contribution-based method is potentially advantageous over a gradient-based method as it does not require the differentiability of a prediction model.

Shapley value and its extension The Shapley value, introduced as a fair division method from economics [Shapley, 1953a], has been deployed in various ML problems. One leading application is data valuation, where the main goal is to quantify the impact of individual data points in model training. Ghorbani and Zou [2019] and Jia et al. [2019] propose to use the Shapley value for measuring the individual data value, and this concept has been extended to handle the randomness of data [Ghorbani et al., 2020, Kwon et al., 2021]. As for the other applications of the Shapley value, model explainability [Ghorbani and Zou, 2020], model valuation [Rozemberczki and Sarkar, 2021], federated learning [Liu et al., 2021], and multi-agent reinforcement learning [Li et al., 2021] have been studied. We refer to Rozemberczki et al. [2022] for a complementary literature review of ML applications of the Shapley value.

The relaxation of the Shapley axioms has been one of the central topics in cooperative game theory [Shapley, 1953b, Banzhaf III, 1964, Kalai and Samet, 1987, Weber, 1988]. Recently, Kwon and Zou [2021] propose to relax the efficiency axiom in the data valuation problem, showing promising results in the low-quality data detection task. Given that Shapley axioms are often not readily applicable to ML problems, relaxing them has the potential to capture a better notion of significance. In this work, we explore the benefits of relaxation of the efficiency axiom on the attribution problem.

2 Preliminaries

We review the marginal contribution and the Shapley value in the context of an attribution problem. We first introduce some notations. For $d \in \mathbb{N}$, let $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{Y} \subseteq \mathbb{R}$ be an input space and an output space, respectively. We use a capital letter $X = (X_1, \dots, X_d)$ for an input random

variable defined on \mathcal{X} , and a lower case letter $x = (x_1, \dots, x_d)$ for its realized value. We denote a prediction model by $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$. For $j \in \mathbb{N}$, we set $[j] := \{1, \dots, j\}$ and denote a power set of $[j]$ by $2^{[j]}$. For a vector $u \in \mathbb{R}^d$ and a subset $S = (j_1, \dots, j_{|S|}) \subseteq [d]$, we denote a subvector by $u_S := (u_{j_1}, \dots, u_{j_{|S|}})$. We assume that X has a joint distribution $p(X)$ such that a conditional distribution $p(X_{[d] \setminus S} | X_S)$ is well-defined for any subset $S \subsetneq [d]$. With the notations, a conditional coalition function $v_{x, \hat{f}}^{(\text{cond})} : 2^{[d]} \rightarrow \mathbb{R}$ is defined as follows [Lundberg and Lee, 2017].

$$v_{x, \hat{f}}^{(\text{cond})}(S) := \mathbb{E}[\hat{f}(x_S, X_{[d] \setminus S}) | X_S = x_S] - \mathbb{E}[\hat{f}(X)], \quad (1)$$

where the first expectation is taken with a conditional distribution $p(X_{[d] \setminus S} | X_S = x_S)$ and the second expectation is taken with a joint distribution $p(X)$. Here, we use a slight abuse of notation for $\hat{f}(x_S, X_{[d] \setminus S})$ to describe $f(u)$ where $u_i = x_i$ if $i \in S$, and $u_i = X_i$ otherwise. By convention, we set $v_{x, \hat{f}}^{(\text{cond})}([d]) := \hat{f}(x) - \mathbb{E}[\hat{f}(X)]$ and $v_{x, \hat{f}}^{(\text{cond})}(\emptyset) := 0$. A conditional coalition function defined in Equation (1) is a prediction recovered after observing partial information x_S compared to the null information. For instance, if $S = \emptyset$, the first term becomes the marginal expectation $\mathbb{E}[\hat{f}(X)]$ and nothing is recovered by $S = \emptyset$. For ease of notation, we write $v_{x, \hat{f}}^{(\text{cond})}(S) = v^{(\text{cond})}(S)$ for remaining part of the paper.

Marginal contribution-based attribution methods Given that the goal of the attribution problem is to assign the significance of an individual feature x_i on the prediction $\hat{f}(x)$, its primary challenge is how to measure the influence of the feature x_i . A leading approach is to quantify the difference in the conditional coalition function values $v^{(\text{cond})}$ after adding one feature of interest. We formalize this concept below.

Definition 1 (Marginal contribution). *For $i, j \in [d]$, we define the marginal contribution of the i -th feature x_i with respect to $j - 1$ features as follows.*

$$\Delta_j(x_i) := \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} v^{(\text{cond})}(S \cup \{i\}) - v^{(\text{cond})}(S). \quad (2)$$

The marginal contribution $\Delta_j(x_i)$ considers every possible subset $S \subseteq [d] \setminus \{i\}$ with the coalition size $|S| = j - 1$ and takes a simple average of the difference $v^{(\text{cond})}(S \cup \{i\}) - v^{(\text{cond})}(S)$. That is, it measures the average contribution of the i -th feature x_i when it is added to a subset S .

Different marginal contributions $\Delta_j(x_i)$ have been studied depending on the coalition size j in the literature. Zintgraf et al. [2017] considered $j = d$ and measured the leave-one-out marginal contribution $\Delta_d(x_i) = \hat{f}(x) - \mathbb{E}[\hat{f}(x_{[d] \setminus \{i\}}, X_i) | X_{[d] \setminus \{i\}} = x_{[d] \setminus \{i\}}]$ as an influence of a feature. Guyon and Elisseeff [2003] considered $j = 1$ and measured the coefficient of determination as an influence of a feature. Although they did not use an individual prediction, their idea is essentially similar to using $\Delta_1(x_i) = \mathbb{E}[\hat{f}(x_i, X_{[d] \setminus \{i\}}) | X_i = x_i] - \mathbb{E}[\hat{f}(X)]$.

Another widely used marginal contribution-based method is the Shapley value [Lundberg and Lee, 2017, Covert and Lee, 2021]. It summarizes the impact of one feature by taking a simple average across all marginal contributions. To be more specific, the Shapley value is defined as follows.

$$\phi_{\text{shap}}(x_i) := \frac{1}{d} \sum_{j=1}^d \Delta_j(x_i). \quad (3)$$

The Shapley value in (3) is known as the unique function that satisfies the four axioms of a fair division in cooperative game theory [Shapley, 1953a]. The four axioms and the uniqueness of the Shapley value are discussed in more detail in Appendix.

Although the Shapley value provides a principled framework in game theory, one critical issue is that the economic notion of the Shapley axioms is not intuitively applicable to the attribution problem [Kumar et al., 2020, Rozemberczki et al., 2022]. In particular, the efficiency axiom, which requires the sum of the attributions to be equal to $v^{(\text{cond})}([d])$, is not necessarily essential because an order of attributions is invariant to the constant multiplication. For instance, for any positive constant $C > 0$,

an attribution $\phi_C(x_i) := C \times \phi_{\text{shap}}(x_i)$ will have the same order as the Shapley value ϕ_{shap} , but the efficiency axiom is not required for ϕ_C . In Section 4, we will revisit this point and introduce a new attribution method that relaxes the efficiency axiom.

Evaluation metrics for the attribution problem. In the literature, different notions of goodness have been proposed, for instance, the complete axiom [Sundararajan et al., 2017, Shrikumar et al., 2017], the local Lipschitzness [Alvarez-Melis and Jaakkola, 2018], and the explanation infidelity [Yeh et al., 2019] with a focus on the total sum of attributions or the sensitivity of attributions. Recently, Jethani et al. [2021a] suggested using the *Inclusion AUC* to assess the goodness of an order of attributions. Specifically, the Inclusion AUC is measured as follows: Given an attribution method, features are first ranked based on their attribution values. Then the area under the receiver operating characteristic curve (AUC) is iteratively evaluated by adding features one by one from the most influential to the least influential. This procedure generates a AUC curve as a function of the number features added, and the Inclusion AUC is defined as the area under this curve. Similar evaluation metrics have been used in Petsiuk et al. [2018] and Lundberg et al. [2020]. Following the literature, we consider the area under the prediction recovery error curve (AUP) defined as follows.

Definition 2 (Area under the prediction recovery error curve). *For a given attribution method ϕ , an input $x \in \mathcal{X}$, and $k \in [d]$, let $\mathcal{I}(k; \phi, x) \subseteq [d]$ be a set of k integers that indicates k most influential features based on their absolute value $|\phi(x_j)|$. For a prediction model \hat{f} , we define the area under the prediction recovery error curve at x as follows.*

$$\text{AUP}(\phi; x, \hat{f}) := \sum_{k=1}^d \left| \hat{f}(x) - \mathbb{E}[\hat{f}(X) \mid X_{\mathcal{I}(k; \phi, x)} = x_{\mathcal{I}(k; \phi, x)}] \right|. \quad (4)$$

AUP is defined as the sum of the absolute differences between the original prediction $\hat{f}(x)$ and its conditional expectation $\mathbb{E}[\hat{f}(X) \mid X_{\mathcal{I}(k; \phi, x)} = x_{\mathcal{I}(k; \phi, x)}]$ when the k most influential features are given. Each term in Equation (4) measures the amount of a prediction that is not recovered by the k most influential features, and thus this prediction recovery error is expected to decrease as k increases. The prediction recovery error can be described as a function of k , and the AUP measures the area under this function as in the Inclusion AUC.

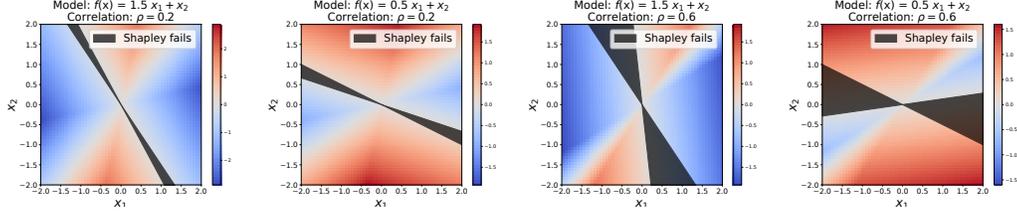
3 The Shapley value is suboptimal

In this section, we show that the suboptimality of the Shapley value through a rigorous analysis of the marginal contribution. We first derive a useful closed-form expression of the marginal contribution when \hat{f} is linear and $p(X)$ is Gaussian (Section 3.1). With this theoretical result, we present two simulation experiments where the Shapley value incorrectly reflects the influence of features, resulting in a suboptimal order of attributions (Section 3.2).

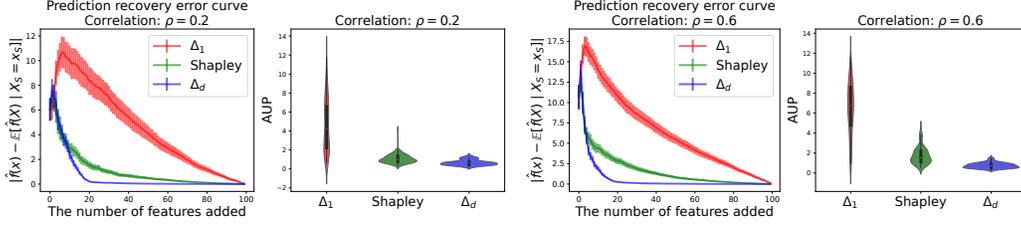
3.1 A closed-form expression of the marginal contribution

To this end, we assume that a prediction model \hat{f} is linear and an input distribution $p(X)$ follows a Gaussian distribution with zero mean and a block diagonal covariance matrix. To be more specific, we define some notations. For $B \in \mathbb{N}$, we set a vector $\mathbf{d} = (d_1, \dots, d_B) \in \mathbb{N}^B$ such that $\sum_{b=1}^B d_b = d$ and a vector $\boldsymbol{\rho} := (\rho_1, \dots, \rho_B) \in [0, 1]^B$. We denote a $d \times d$ block diagonal covariance matrix by $\Sigma_{\boldsymbol{\rho}, \mathbf{d}}^{(\text{block})} = \text{diag} \left(\Sigma_{(\rho_1, d_1)}, \dots, \Sigma_{(\rho_B, d_B)} \right)$ where $\Sigma_{(\rho_b, d_b)} = (1 - \rho_b)I_{d_b} + \rho_b \mathbf{1}_{d_b} \mathbf{1}_{d_b}^T$. Here, for $j \in \mathbb{N}$, we denote the $j \times j$ identity matrix by I_j , the j -dimensional vector of ones by $\mathbf{1}_j := (1, \dots, 1)^T \in \mathbb{R}^j$ and $\mathbf{0}_j := 0 \times \mathbf{1}_j$. Lastly, we denote a Gaussian distribution with a mean vector μ and a covariance matrix Σ by $\mathcal{N}(\mu, \Sigma)$. With the notations, we assume $X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{\boldsymbol{\rho}, \mathbf{d}}^{(\text{block})})$. That is, every feature is normalized to have a unit variance and is included in one of B independent clusters. For $j \in [B]$, the size of the j -th cluster is d_j , and features are equally correlated to each other within a cluster. The correlation levels can vary from cluster to cluster.

In general, the marginal contribution in Equation (2) does not have a closed-form expression, and it makes a rigorous analysis of the Shapley value difficult. In the following theorem, we derive a closed-form expression of the marginal contribution when \hat{f} is linear and $X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{\boldsymbol{\rho}, \mathbf{d}}^{(\text{block})})$.



(a) Illustrations of the suboptimality of Shapley-based feature attributions on the four different situations when $d = 2$.



(b) Illustrations of a prediction recovery error curve and AUP comparison when $d = 100$.

Figure 1: The Shapley value is suboptimal. (Top) a region described in black denotes the area the Shapley value fails to select the more influential feature. We encode $\mathcal{E}(1; x, \hat{f}) - \mathcal{E}(2; x, \hat{f})$ as the background color to visualize which feature is more influential. Blue color describes a region where the first feature x_1 is more influential, *i.e.*, $\mathcal{E}(1; x, \hat{f}) < \mathcal{E}(2; x, \hat{f})$, and red describes a region where the second feature x_2 is more influential, *i.e.*, $\mathcal{E}(1; x, \hat{f}) > \mathcal{E}(2; x, \hat{f})$. The intensity for $\mathcal{E}(1; x, \hat{f}) - \mathcal{E}(2; x, \hat{f})$ is described in a color bar. (Bottom) we compare the three attribution methods Δ_1 , ϕ_{shap} , and Δ_d on the two different situations by varying correlation $\rho \in \{0.2, 0.6\}$. As for the prediction recovery error curve, we denote a 95% confidence band based on 100 samples. The lower AUP is, the better. In both settings, the Shapley value is suboptimal according to AUP.

Theorem 1 (A closed-form expression for the marginal contribution). *Suppose $\hat{f}(x) = \hat{\beta}_0 + x^T \hat{\beta}$ for some $(\hat{\beta}_0, \hat{\beta}) \in \mathbb{R} \times \mathbb{R}^d$ and $X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{\rho, d}^{(\text{block})})$. Then, for $i, j \in [d]$, the marginal contribution of the i -th feature x_i with respect to $j - 1$ samples is expressed as*

$$\Delta_j(x_i) = x^T H(i, j) \hat{\beta},$$

for some explicit matrix $H(i, j) \in \mathbb{R}^{d \times d}$.

A proof and the explicit term for $H(i, j)$ are provided in Appendix. Theorem 1 shows that the marginal contribution is a bilinear function of an input x and the estimated regression coefficient $\hat{\beta}$. One direct consequence is that the Shapley value also has a bilinear form $\phi_{\text{shap}}(x_i) = x^T H(i) \hat{\beta}$ for $H(i) := \sum_{j=1}^d H(i, j) / d$. We emphasize that this bilinear form greatly improves computational efficiency. Specifically, for all $i, j \in [d]$, since the term $H(i, j) \hat{\beta}$ is not a function of an input x , we only need to compute the one-time in multiple attribution computations. Moreover, it also leads to a memory efficient algorithm as there is no need to store the $d \times d$ matrix $H(i, j)$.

3.2 Motivational examples

With the theoretical result introduced in the previous subsection, we show that the Shapley-based feature attribution is suboptimal and fails to assign larger values to more influential features.

When there are two features. When $d = 2$, there are only two possible values for AUP. For any attribution method ϕ ,

$$\text{AUP}(\phi; x, \hat{f}) = \begin{cases} \mathcal{E}(1; x, \hat{f}) & \text{if } \mathcal{I}(1; \phi, x) = \{1\} \\ \mathcal{E}(2; x, \hat{f}) & \text{otherwise} \end{cases},$$

where $\mathcal{E}(k; x, \hat{f}) := \left| \hat{f}(x_1, x_2) - \mathbb{E}[\hat{f}(X_1, X_2) \mid X_k = x_k] \right|$ for $k \in \{1, 2\}$. Therefore, the optimal order based on AUP is fully determined by $\mathcal{E}(1; x, \hat{f})$ and $\mathcal{E}(2; x, \hat{f})$, for instance, the first feature x_1 is more influential than the second one x_2 if $\mathcal{E}(1; x, \hat{f}) < \mathcal{E}(2; x, \hat{f})$. It is intuitively sensible because $\mathcal{E}(1; x, \hat{f}) < \mathcal{E}(2; x, \hat{f})$ means that the original prediction $\hat{f}(x)$ is more accurately recovered by the first feature x_1 than the second one x_2 .

Using the optimal order, we demonstrate that the Shapley value does not necessarily assign a large attribution to a more influential feature. We consider the four different scenarios with two different prediction models $\hat{f}(x) \in \{1.5x_1 + x_2, 0.5x_1 + x_2\}$ and two different Gaussian distributions, $X \sim \mathcal{N}(\mathbf{0}_2, \Sigma_{(\rho, 2)})$ for $\rho \in \{0.2, 0.6\}$. In these four scenarios, the terms $\mathcal{E}(1; x, \hat{f})$ and $\mathcal{E}(2; x, \hat{f})$ have a closed-form expression, and thus the optimal order is explicitly obtained. Moreover, due to Theorem 1, a more influential feature according to the Shapley value is explicitly obtained.

Figure 1(a) illustrates the suboptimality of the Shapley value on the four different situations. In any situation, there is a non-negligible region (described in black) where the Shapley value fails to select a more influential feature. In addition, this suboptimal area increases as the correlation gets larger, showing that the Shapley value-based explainability becomes poor when features are highly correlated.

When there are more than two features When $d > 2$, it is infeasible to find the exact optimal order because there are 2^{d-1} possible AUPs. For this reason, we compare the Shapley value with the two commonly used marginal contribution-based methods Δ_1 and Δ_d , showing the Shapley value is not optimal in terms of AUP. We assume the following setting: a trained model is linear $\hat{f}(x) = \hat{\beta}_0 + x^T \hat{\beta}$ for some $(\hat{\beta}_0, \hat{\beta}) \in \mathbb{R} \times \mathbb{R}^d$ and an input vector $X = (X_1, \dots, X_d)$ follows a Gaussian distribution $\mathcal{N}(\mathbf{0}_d, \Sigma_{(\rho, d)})$. That is, there are d features and they are equally correlated to each other with the correlation ρ . We set $d = 100$ and consider two different situations by varying $\rho \in \{0.2, 0.6\}$. Similar to the previous analysis, due to Theorem 1, the three attribution methods are explicitly obtained. We evaluate the prediction recovery error and the AUP on the 100 held-out test samples randomly drawn from the distribution $\mathcal{N}(\mathbf{0}_d, \Sigma_{(\rho, d)})$. Detailed information is provided in Appendix.

Figure 1(b) illustrates the suboptimality of the Shapley value when $d = 100$ and $\rho \in \{0.2, 0.6\}$. In any situation, the prediction recovery curves for the Shapley value (described in green) have a steeper slope than Δ_1 (described in red), but is not optimal as the Δ_d (described in blue) approaches to zero faster. When $\rho = 0.6$, the suboptimality becomes more severe in that the gap between Δ_d and the Shapley value gets larger.

4 Proposed method: WeightedSHAP

Our motivational examples in the previous section suggest that the Shapley value does not necessarily assign larger attributions for more influential features, leading to a suboptimal order of features. In fact, the last marginal contribution Δ_d outperforms other attribution methods in Figure 1(b). Although the use of Δ_d is promising, we show that focusing on one marginal contribution might lead to an unstable attribution method, suggesting a weighted mean of the marginal contributions. To be more concrete, we first examine the estimation error of the marginal contribution in the following.

Analysis of estimation error In practice, the Shapley value needs to be estimated, resulting in an estimation error. Given the mathematical form of the Shapley value in (3), this estimation error arises from the estimation error of the marginal contribution. In this reason, we investigate the estimation error of the marginal contribution. We consider the same setting used in Figure 1(b) with $\rho = 0.6$. As for the estimation of the marginal contribution, we follow a standard algorithm to estimate a conditional coalition function $v^{(\text{cond})}$ used in Jethani et al. [2021a] and a sampling-based algorithm to approximate Δ_j . A detailed explanation for the estimation procedure and the additional result for $\rho = 0.2$ are provided in Appendix.

Figure 2 shows the relative difference between the true marginal contribution Δ_j obtained by Theorem 1 and its estimate $\hat{\Delta}_j$ as a function of the coalition size $j \in [d]$. Here, we use the relative

difference between A and B defined as $|A - B| / \max(|A|, |B|)$ to avoid numerical instability that can be occurred by too small marginal contribution values. It shows the Δ_d , the most informative marginal contribution in Figure 1(b), has the largest relative difference from the true value. In other words, Δ_d has the largest signal to explain a model prediction, but at the same time, it is the most unstable in terms of the estimation error. This finding motivates us to consider a weighted mean of the marginal contributions that can reduce the estimation error while maintaining signals.

Proposed method For a weighted vector $\mathbf{w} = (w_1, \dots, w_d)$ such that $\sum_{j=1}^d w_j = 1$ and $w_j \geq 0$ for all $j \in [d]$, we consider a weighted mean of the marginal contributions

$$\phi_{\mathbf{w}}(x_i) := \sum_{j=1}^d w_j \Delta_j(x_i). \quad (5)$$

A weighted mean $\phi_{\mathbf{w}}(x_i)$ is expected to capture the influence of features better than the Shapely value (3) by assigning a large weight to important marginal contributions. As for the game theoretic interpretation, a mathematical form of Equation (5) is known as a semivalue in cooperative game theory. It satisfies all the Shapley axioms but the efficiency axiom, which is not crucial in the attribution as we discussed in Section 2 [Dubey and Weber, 1977, Ridaoui et al., 2018]. Due to the relaxation of the efficiency axiom, a semivalue is not uniquely determined, but it is known that the semivalue is *almost* unique up to a weighted mean operation. A detailed explanation for the semivalue is provided in Appendix.

One natural and practical question that arises when using a weighted mean $\phi_{\mathbf{w}}(x_i)$ is how to select the weight vector \mathbf{w} . Since a weight vector to be selected is desired to have a certain good property, this question can be rephrased as to which concept of goodness should be optimized. However, it is difficult to have one universal desideratum by the intricate nature of model interpretations. There are different types of goodness and they often represent independent characteristics, as we discussed in Section 2. In other words, a good attribution essentially depends on a practitioner’s downstream task. To reflect this, we propose to learn a weight vector that optimizes a user-defined utility. To be more specific, we let $\mathcal{W} \subseteq \{w \in \mathbb{R}^d : \sum_{j=1}^d w_j = 1, w_j \geq 0\}$ be a parametrized family of weights and \mathcal{T} be a user-defined utility function that takes as input an attribution method and outputs its utility. Without loss of generality, we assume that the larger \mathcal{T} is, the better it is (e.g., the negative value of AUP). Given \mathcal{T} and \mathcal{W} , we propose WeightedSHAP as follows.

$$\phi^{\text{WeightedSHAP}}(\mathcal{T}, \mathcal{W}) := \phi_{\mathbf{w}^*(\mathcal{T}, \mathcal{W})}, \quad (6)$$

where $\mathbf{w}^*(\mathcal{T}, \mathcal{W}) := \operatorname{argmax}_{\mathbf{w} \in \mathcal{W}} \mathcal{T}(\phi_{\mathbf{w}})$. That is, we learn the optimal weight by optimizing a user-defined utility.

When \mathcal{W} includes the uniform weight $(1/d, \dots, 1/d) \in \mathbb{R}^d$, then by its construction, we can guarantee that WeightedSHAP is always better than or equal to the Shapley value according to the utility \mathcal{T} . For instance, when the negative value of AUP is used for the utility \mathcal{T} , AUP of WeightedSHAP is less than that of the Shapley value, i.e., $\text{AUP}(\phi^{\text{WeightedSHAP}}) \leq \text{AUP}(\phi_{\text{shap}})$. Moreover, the more weight vectors are in \mathcal{W} , the better the quality of $\phi^{\text{WeightedSHAP}}$ is guaranteed. WeightedSHAP $\phi^{\text{WeightedSHAP}}$ depends on a set \mathcal{W} . In our experiments, we parameterize an element $\mathbf{w} \in \mathcal{W}$ by the Beta distribution inspired by mathematical properties of the semivalue in Monderer and Samet [2002, Theorem 11]. Detailed information is provided in Appendix.

Example 1 (WeightedSHAP and the Shapley value ϕ_{shap} on AUP). *We revisit the motivational example introduced in Section 3.2. With the negative AUP for \mathcal{T} and some $\mathcal{W} \supseteq \{\Delta_d, \phi_{\text{shap}}\}$, WeightedSHAP achieves significantly lower AUP than both Δ_d and ϕ_{shap} . Specifically, when $(d, \rho) = (100, 0.6)$, the AUPs of $(\Delta_d, \phi_{\text{shap}}, \phi^{\text{WeightedSHAP}})$ are $(1.49 \pm 0.06, 1.65 \pm 0.08, 0.77 \pm 0.03)$, respectively, where the numbers denote “mean \pm standard error” based on the 100 held-out test*

Relative difference btw Δ_j and $\hat{\Delta}_j$
Correlation: $\rho = 0.6$

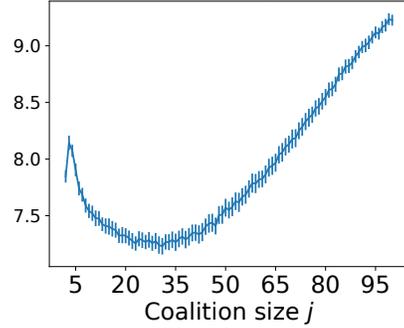
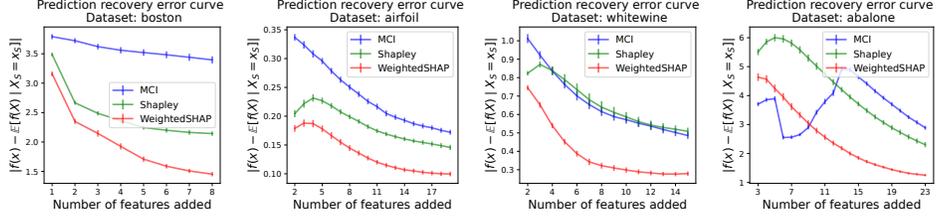
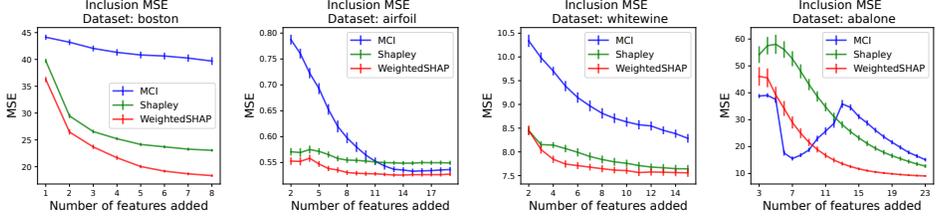


Figure 2: Illustrations of the relative difference between the true marginal contribution Δ_j and its estimate $\hat{\Delta}_j$ as a function of the coalition size $j \in [d]$. We consider the same setting used in Figure 1(b). The Δ_d is shown to have the largest relative difference.



(a) Illustrations of the prediction recovery error curve on the four regression datasets. The lower, the better.



(b) Illustrations of the Inclusion MSE curve on the four regression datasets. The lower, the better.

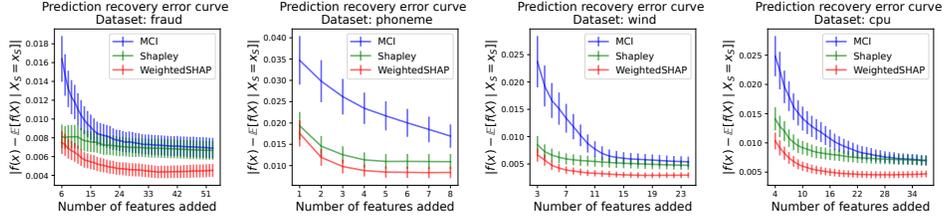
Figure 3: Regression tasks. Illustrations of the prediction recovery error curve and the Inclusion MSE curve as a function of the number of features added. We add features from most influential to the least influential. We denote a 95% confidence interval based on 50 independent runs. WeightedSHAP achieves a significantly smaller MSE with fewer features than the MCI and the Shapley value.

samples. Meanwhile, the estimation errors of $(\Delta_d, \phi_{\text{shap}}, \phi_{\text{WeightedSHAP}})$ are $(9.23 \pm 0.02, 6.16 \pm 0.04, 7.90 \pm 0.11)$, respectively. In short, WeightedSHAP achieves a significantly lower estimation error than Δ_d while achieving the lowest AUP. Although ϕ_{shap} achieves the lowest estimation error, its AUP is significantly greater than both Δ_d and $\phi_{\text{WeightedSHAP}}$. The uniform weight in ϕ_{shap} helps reduce the estimation error, but it loses signals too much. In contrast, WeightedSHAP well balances the signal and the estimation error, i.e., reducing the estimation error while taking more signals.

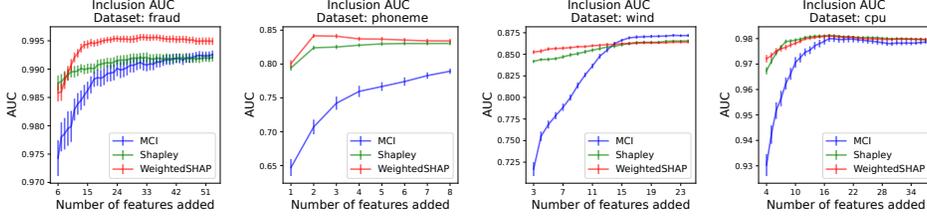
Implementation algorithm for WeightedSHAP Given a finite set \mathcal{W} and an easy-to-compute utility function \mathcal{T} , the optimal weight \mathbf{w}^* can be achieved by iteratively evaluating the utility \mathcal{T} for each attribution method $\phi_{\mathbf{w}}$ with $\mathbf{w} \in \mathcal{W}$. In addition, $\phi_{\mathbf{w}}$ is readily obtained as long as there are the marginal contribution estimates. Therefore, the key part of the implementation algorithm is to estimate a set of marginal contributions. The estimation of the marginal contributions consists of two parts, estimation of a conditional coalition function $v^{(\text{cond})}$ and approximation of the marginal contribution Δ_j . As for the first part, we train a surrogate model that takes as input a subset of input features and outputs a conditional expectation of a prediction value given the same subset [Frye et al., 2020, Jethani et al., 2021a,b]. It is known that this surrogate model unbiasedly estimates a conditional expectation of a prediction value given a subset of features under mild conditions [Frye et al., 2020, Covert et al., 2020]. Regarding the second part, a weighted mean is approximated by a sampling-based algorithm [Ghorbani and Zou, 2019, Kwon and Zou, 2021]. We provide a pseudo algorithm in Appendix. In terms of the computational cost, our algorithm is comparable to a standard the Shapley value estimation algorithm because both algorithms need to estimate the marginal contributions as a primary part [Lundberg and Lee, 2017, Frye et al., 2020]. For instance, with the classification dataset fraud, the marginal contribution estimation part takes 20.7 seconds per sample on average but the weight optimization part only takes 0.18 seconds, i.e., the weight optimization part is only 0.86% of the total compute.

5 Experimental results

We demonstrate the practical efficacy of WeightedSHAP on various regression and classification datasets. We compare WeightedSHAP with the marginal contribution importance (MCI) proposed by Catav et al. [2021] and the Shapley value on the prediction recovery error task and the inclusion performance task. Each task assesses the goodness of an attribution order by iteratively



(a) Illustrations of the prediction recovery error curve on the four binary classification datasets. The lower, the better.



(b) Illustrations of the Inclusion AUC curve on the four binary classification datasets. The higher, the better.

Figure 4: **Classification tasks.** Illustrations of the prediction recovery error curve and the Inclusion AUC curve as a function of the number of features added. Details are provided in Figure 3. WeightedSHAP achieves a significantly higher AUC with fewer features than the MCI and the Shapley value.

measuring how much the original model prediction or its performance is recovered with a given number of features. As for the model performance, we evaluate mean squared error (MSE) and AUC for regression and classification problems, respectively. We consider a gradient boosting model for a prediction model \hat{f} . As for the surrogate model in coalition function estimation $v^{(\text{cond})}$, we use a multilayer perceptron model with the two hidden layers, and each layer has 128 neurons and the ELU activation function [Clevert et al., 2015]. As for the WeightedSHAP in (6), we use the negative value of the AUP for \mathcal{T} and a set \mathcal{W} that has 13 different weights including Δ_1 , ϕ_{shap} , and Δ_d . All the missing details about numerical experiments are provided in Appendix, and our Python-based implementations are available at <https://github.com/ykwon0407/WeightedSHAP>.

Figure 3(a) compares the prediction recovery error curves for the WeightedSHAP (described in red) with the MCI (described in blue) and the Shapley value (described in green). WeightedSHAP shows always lower prediction recovery errors than the MCI and the Shapley value. Given that WeightedSHAP minimizes the AUP, which is the sum of prediction recovery error $|\hat{f}(x) - \mathbb{E}[\hat{f}(X) | X_S = x_S]|$, WeightedSHAP does not necessarily have a smaller prediction recovery error for every number of features added. As for the MSE, Figure 3(b) shows that WeightedSHAP has a significantly smaller MSE than baseline methods with fewer features. In particular, on the `airfoil` dataset, WeightedSHAP achieves 0.53 MSE with 10 features, but the Shapley value never achieves this value because of the suboptimality of the attribution order.

We also evaluate the prediction recovery error and AUC for the four classification datasets. Similar to the regression cases, Figures 4(a) and 4(b) show that WeightedSHAP effectively assigns larger values for more influential features and recovers the original prediction $\hat{f}(x)$ significantly faster than the MCI and the Shapley value. Specifically, on the `fraud` dataset, WeightedSHAP achieves 0.995 AUC with 14 features, but the Shapley value always has the lower AUC value. Our findings are consistently observed with a different model for \hat{f} or other datasets. Additional experimental results with different evaluation metrics and a qualitative assessment of WeightedSHAP are provided in Appendix.

5.1 Illustrative examples from MNIST

We present a qualitative assessment of WeightedSHAP and examine how its top influential features differ from those from SHAP using the MNIST dataset. We train a convolutional neural network model using the same setting suggested in Jethani et al. [2021a]. It achieves 98.6 % accuracy

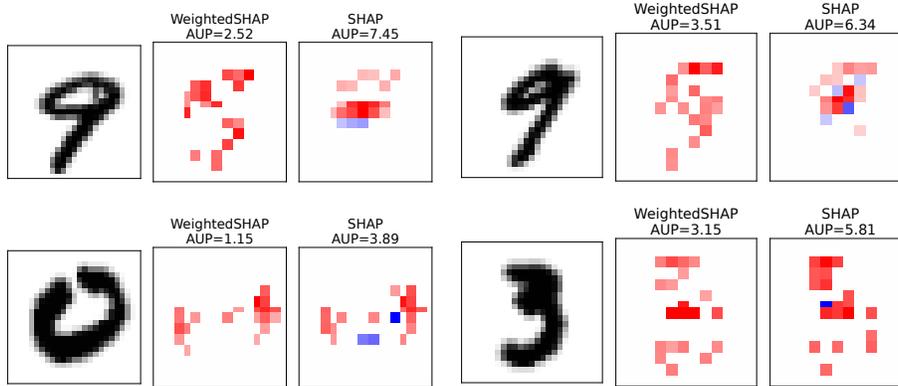


Figure 5: Illustrative examples of WeightedSHAP and SHAP attributions on MNIST images. We present the top 10% of influential features. Red (*resp.* blue) color indicates the corresponding feature positively (*resp.* negatively) affects the model prediction. (Top) WeightedSHAP clearly captures the last stroke of nine while SHAP fails to capture it. (Bottom) While SHAP has noisy negative feature attributions described by blue pixels, WeightedSHAP provides noiseless and intuitive explanations.

on the test dataset. We select illustrative images with a significant difference in AUPs between WeightedSHAP and SHAP.

Figure 5 compares the top 10% influential attributions for WeightedSHAP and the Shapley value. On the top images, while SHAP fails to capture the last stroke of digit nine, which is a crucially important stroke to differentiate from the digit zero, WeightedSHAP clearly captures the strokes. On the bottom images, SHAP produces unintuitive negative attributions, providing noisy explanations. In contrast, WeightedSHAP provides noiseless and intuitive explanations.

6 Conclusion

In this paper, we provide an analysis of the widely used SHAP attribution method. We discover that even in simple natural settings, SHAP can incorrectly identify important features. Mathematically, a key limitation of SHAP is in that it assigns uniform weights to all marginal contributions. We propose WeightedSHAP which generalizes the Shapley value by relaxing the efficiency axiom. WeightedSHAP learns to pay more attention to the marginal contributions that have more signal on a prediction, assigning larger attributions for more influential features. There are several limitations of WeightedSHAP that motivate interesting future works. Here we use the AUP metric to optimize the weights because AUP is commonly used in practice. However, there is no agreed-upon metric for evaluating feature attribution methods. Different users may care about different notions of attribution. Developing variants of marginal contribution weighting optimized for different applications is an important direction of future research. We believe that the core contribution of this paper—that the uniform weighting used by SHAP can be suboptimal—still provides useful insights for these investigations. Here we focus our experiments on directly comparing WeightedSHAP with SHAP because our goal is to characterize the limitations of SHAP. There is a large body of works comparing SHAP with other attribution methods that are complementary to our work [Yeh et al., 2019, Jethani et al., 2021a].

Acknowledgment

The authors would like to thank all anonymous reviewers for their constructive comments. We also would like to thank Ruishan Liu for the helpful discussion.

References

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Proceedings of the 31st international conference on neural information processing systems*, pages

- 4768–4777, 2017.
- Dominik Janzing, Lenon Minorics, and Patrick Blöbaum. Feature relevance quantification in explainable ai: A causal problem. In *International Conference on Artificial Intelligence and Statistics*, pages 2907–2916. PMLR, 2020.
- Mukund Sundararajan and Amir Najmi. The many shapley values for model explanation. In *International conference on machine learning*, pages 9269–9278. PMLR, 2020.
- Christopher Frye, Damien de Mijolla, Tom Begley, Laurence Cowton, Megan Stanley, and Ilya Feige. Shapley explainability on the data manifold. *arXiv preprint arXiv:2006.01272*, 2020.
- Kjersti Aas, Martin Jullum, and Anders Løland. Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *Artificial Intelligence*, 298:103502, 2021.
- Tom Heskes, Evi Sijben, Ioan Gabriel Bucur, and Tom Claassen. Causal shapley values: Exploiting causal knowledge to explain individual predictions of complex models. *Advances in neural information processing systems*, 33:4778–4789, 2020.
- Jiaxuan Wang, Jenna Wiens, and Scott Lundberg. Shapley flow: A graph-based approach to interpreting model predictions. In *International Conference on Artificial Intelligence and Statistics*, pages 721–729. PMLR, 2021.
- 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.
- 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.
- Neil Jethani, Mukund Sudarshan, Ian Covert, Su-In Lee, and Rajesh Ranganath. Fastshap: Real-time shapley value estimation. *arXiv preprint arXiv:2107.07436*, 2021a.
- Scott M Lundberg, Bala Nair, Monica S Vavilala, Mayumi Horibe, Michael J Eisses, Trevor Adams, David E Liston, Daniel King-Wai Low, Shu-Fang Newman, Jerry Kim, et al. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature biomedical engineering*, 2(10):749–760, 2018.
- Joseph D Janizek, Ayse Berceste Dincer, Safiye Celik, Hugh Chen, William Chen, Kamila Naxerova, and Su-In Lee. Uncovering expression signatures of synergistic drug response using an ensemble of explainable ai models. *bioRxiv*, 2021.
- Wei Qiu, Hugh Chen, Ayse Berceste Dincer, Scott Lundberg, Matt Kaeberlein, and Su-In Lee. Interpretable machine learning prediction of all-cause mortality. *medRxiv*, pages 2021–01, 2022.
- Stan Lipovetsky and Michael Conklin. Analysis of regression in game theory approach. *Applied Stochastic Models in Business and Industry*, 17(4):319–330, 2001.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Art B Owen. Sobol’ indices and shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.
- Baptiste Broto, François Bachoc, and Marine Depecker. Variance reduction for estimation of shapley effects and adaptation to unknown input distribution. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):693–716, 2020.
- Qingyuan Zhao and Trevor Hastie. Causal interpretations of black-box models. *Journal of Business & Economic Statistics*, 39(1):272–281, 2021.
- Clément Bénéard, Gérard Biau, Sébastien Da Veiga, and Erwan Scornet. Shaff: Fast and consistent shapley effect estimates via random forests. In *International Conference on Artificial Intelligence and Statistics*, pages 5563–5582. PMLR, 2022.

- Leo Breiman, Jerome H Friedman, Richard A Olshen, and Charles J Stone. *Classification and regression trees*. Routledge, 2017.
- Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *International Conference on Machine Learning*, pages 883–892. PMLR, 2018.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR, 2017.
- Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. Towards better understanding of gradient-based attribution methods for deep neural networks. *arXiv preprint arXiv:1711.06104*, 2017.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. *Advances in neural information processing systems*, 31, 2018.
- Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953a.
- Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, pages 2242–2251. PMLR, 2019.
- Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nezihe Merve Gurel, Bo Li, Ce Zhang, Costas Spanos, and Dawn Song. Efficient task-specific data valuation for nearest neighbor algorithms. *Proceedings of the VLDB Endowment*, 12(11):1610–1623, 2019.
- Amirata Ghorbani, Michael Kim, and James Zou. A distributional framework for data valuation. In *International Conference on Machine Learning*, pages 3535–3544. PMLR, 2020.
- Yongchan Kwon, Manuel A Rivas, and James Zou. Efficient computation and analysis of distributional shapley values. In *International Conference on Artificial Intelligence and Statistics*, pages 793–801. PMLR, 2021.
- Amirata Ghorbani and James Y Zou. Neuron shapley: Discovering the responsible neurons. *Advances in Neural Information Processing Systems*, 33:5922–5932, 2020.
- Benedek Rozemberczki and Rik Sarkar. The shapley value of classifiers in ensemble games. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1558–1567, 2021.
- Zelei Liu, Yuanyuan Chen, Han Yu, Yang Liu, and Lizhen Cui. Gtg-shapley: Efficient and accurate participant contribution evaluation in federated learning. *arXiv preprint arXiv:2109.02053*, 2021.
- Jiahui Li, Kun Kuang, Baoxiang Wang, Furui Liu, Long Chen, Fei Wu, and Jun Xiao. Shapley counterfactual credits for multi-agent reinforcement learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 934–942, 2021.
- Benedek Rozemberczki, Lauren Watson, Péter Bayer, Hao-Tsung Yang, Olivér Kiss, Sebastian Nilsson, and Rik Sarkar. The shapley value in machine learning. *arXiv preprint arXiv:2202.05594*, 2022.
- Lloyd S Shapley. *Additive and non-additive set functions*. Princeton University, 1953b.
- John F Banzhaf III. Weighted voting doesn’t work: A mathematical analysis. *Rutgers L. Rev.*, 19: 317, 1964.

- Ehud Kalai and Dov Samet. On weighted shapley values. *International journal of game theory*, 16(3):205–222, 1987.
- Robert J Weber. Probabilistic values for games. *The Shapley Value. Essays in Honor of Lloyd S. Shapley*, pages 101–119, 1988.
- Yongchan Kwon and James Zou. Beta shapley: a unified and noise-reduced data valuation framework for machine learning. *arXiv preprint arXiv:2110.14049*, 2021.
- Luisa M Zintgraf, Taco S Cohen, Tameem Adel, and Max Welling. Visualizing deep neural network decisions: Prediction difference analysis. *arXiv preprint arXiv:1702.04595*, 2017.
- Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of machine learning research*, 3(Mar):1157–1182, 2003.
- I Elizabeth Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, and Sorelle Friedler. Problems with shapley-value-based explanations as feature importance measures. In *International Conference on Machine Learning*, pages 5491–5500. PMLR, 2020.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International conference on machine learning*, pages 3145–3153. PMLR, 2017.
- David Alvarez-Melis and Tommi S Jaakkola. On the robustness of interpretability methods. *arXiv preprint arXiv:1806.08049*, 2018.
- Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun Suggala, David I Inouye, and Pradeep K Ravikumar. On the (in) fidelity and sensitivity of explanations. *Advances in Neural Information Processing Systems*, 32, 2019.
- Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*, 2018.
- Pradeep Dubey and Robert James Weber. Probabilistic values for games. Technical report, YALE UNIV NEW HAVEN CONN COWLES FOUNDATION FOR RESEARCH IN ECONOMICS, 1977.
- Mustapha Ridaoui, Michel Grabisch, and Christophe Labreuche. An axiomatisation of the banzhaf value and interaction index for multichoice games. In *International Conference on Modeling Decisions for Artificial Intelligence*, pages 143–155. Springer, 2018.
- Dov Monderer and Dov Samet. Variations on the shapley value. *Handbook of game theory with economic applications*, 3:2055–2076, 2002.
- Neil Jethani, Mukund Sudarshan, Yindalon Aphinyanaphongs, and Rajesh Ranganath. Have we learned to explain?: How interpretability methods can learn to encode predictions in their interpretations. In *International Conference on Artificial Intelligence and Statistics*, pages 1459–1467. PMLR, 2021b.
- Ian Covert, Scott M Lundberg, and Su-In Lee. Understanding global feature contributions with additive importance measures. *Advances in Neural Information Processing Systems*, 33:17212–17223, 2020.
- Amnon Catav, Boyang Fu, Yazeed Zoabi, Ahuva Libi Weiss Meilik, Noam Shomron, Jason Ernst, Sriram Sankararaman, and Ran Gilad-Bachrach. Marginal contribution feature importance - an axiomatic approach for explaining data. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1324–1335. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/catav21a.html>.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Dootika Vats and Christina Knudson. Revisiting the gelman–rubin diagnostic. *Statistical Science*, 36(4):518–529, 2021.
- Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- David Harrison Jr and Daniel L Rubinfeld. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 5(1):81–102, 1978.
- Andrea Dal Pozzolo, Olivier Caelen, Reid A Johnson, and Gianluca Bontempi. Calibrating probability with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence*, pages 159–166. IEEE, 2015.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.
- 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*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] We describe the limitations in the conclusion section.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] We provide a proof in Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We provide used codes in the supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We provide implementation details in Appendix.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] It is not the main focus of our paper.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]

- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

In appendix, we provide proofs in Section A, a detailed discussion on the Shapley value and the semivalue in Section B, implementation details in Section C, and additional experimental results in Section D. Our main findings presented in the manuscript are consistently observed with different datasets and prediction models: the influential features identified by WeightedSHAP better recapitulate the model’s predictions compared to the features identified by the Shapley value. A Python-based implementation code is attached as a part of the supplementary material.

A Proofs

In this section, we provide a proof of Theorem 1. We first prove a lemma that shows a closed-form expression of the marginal contribution when a covariance matrix has an exchangeable correlation structure, *i.e.*, $\Sigma_{(\rho,d)} = (1 - \rho)I_d + \rho\mathbb{1}_d\mathbb{1}_d^T$. To begin with, for $d \in \mathbb{N}$ and $S \in [d]$, we set a d -dimensional vector $e_S \in \{0, 1\}^d$ whose i -th element is one if $i \in S$, otherwise zero. We set $e_i := e_{\{i\}}$ for $i \in [d]$.

Lemma 2 (A closed-form expression of the Marginal contribution when a covariance matrix has an exchangeable correlation structure.). *For $d \in \mathbb{N}$, suppose $\hat{f}(x) = \hat{\beta}_0 + x^T \hat{\beta}$ for some $(\hat{\beta}_0, \hat{\beta}) \in \mathbb{R} \times \mathbb{R}^d$ and $X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{(\rho,d)})$ for $\rho \in [0, 1)$. Then for $i, j \in [d]$, the marginal contribution of the i -th feature x_i with respect to $j - 1$ samples is given as*

$$\Delta_j(x_i) = x^T h_d(i, j) \hat{\beta},$$

where

$$\begin{aligned} & h_d(i, j) \\ & := e_i e_i^T + \frac{\rho}{1 + \rho(j-1)} \frac{d-j}{d-1} e_i (\mathbb{1}_d - e_i)^T - \frac{\rho}{1 - \rho + \rho(j-1)} \frac{j-1}{d-1} (\mathbb{1}_d - e_i) e_i^T \\ & + \left(\frac{\rho}{1 + \rho(j-1)} - \frac{\rho}{1 - \rho + \rho(j-1)} \right) \frac{(j-1)(d-j)}{(d-1)(d-2)} \left((\mathbb{1}_d - e_i)(\mathbb{1}_d - e_i)^T - (I_d - e_i e_i^T) \right). \end{aligned}$$

Proof of Lemma 2. Without loss of generality, we set $\hat{\beta}_0 = 0$ as a constant intercept does not affect the marginal contribution. We denote $\Sigma_{S,S} := (1 - \rho)I_{|S|} + \rho\mathbb{1}_{|S|}\mathbb{1}_{|S|}^T$. Then, for any S , we have

$$\begin{aligned} v^{(\text{cond})}(S) &= x_S^T \hat{\beta}_S + (\Sigma_{D \setminus S, S} \Sigma_{S, S}^{-1} x_S)^T \hat{\beta}_{D \setminus S} \\ &= x_S^T \left(\hat{\beta}_S + \left(\frac{\rho}{1 - \rho} - \frac{|S|\rho^2}{(1 - \rho)(1 + \rho(|S| - 1))} \right) \mathbb{1}_{|S|} \mathbb{1}_{d-|S|}^T \hat{\beta}_{D \setminus S} \right) \\ &= x_S^T \hat{\beta}_S + \frac{\rho}{1 - \rho + \rho|S|} x_S^T \mathbb{1}_{|S|} \mathbb{1}_{d-|S|}^T \hat{\beta}_{D \setminus S} \\ &= x_S^T \hat{\beta}_S + \frac{\rho}{1 - \rho + \rho|S|} x^T e_S e_S^T \hat{\beta}. \end{aligned}$$

Here, the second equality is due to $\Sigma_{S,S}^{-1} = (1 - \rho)^{-1}I_{|S|} - \frac{\rho}{(1-\rho)(1+\rho(|S|-1))} \mathbf{1}_{|S|}\mathbf{1}_{|S|}^T$. Hence, we have

$$v^{(\text{cond})}(S \cup \{i\}) - v^{(\text{cond})}(S) = x_i \hat{\beta}_i - \frac{\rho}{1 - \rho + \rho|S|} x^T e_S e_{S^c}^T \hat{\beta} + \frac{\rho}{1 + \rho|S|} x^T e_{S \cup \{i\}} e_{(S \cup \{i\})^c}^T \hat{\beta}.$$

[Step 1] Since $e_S e_{S^c}^T = e_S (\mathbf{1}_d - e_S)^T$.

$$\begin{aligned} \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T &= \frac{j-1}{d-1} (\mathbf{1}_d - e_i) \mathbf{1}_d^T - \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_S^T \\ &= \frac{j-1}{d-1} (\mathbf{1}_d - e_i) \mathbf{1}_d^T - P(i, j), \end{aligned}$$

where, for $k, l \in [d]$, $P(i, j)_{kl}$ is given by

$$P(i, j)_{kl} = \begin{cases} \frac{\binom{d-2}{j-2}}{\binom{d-1}{j-1}} = \frac{j-1}{d-1} & \text{if } k = l, i \notin \{k, l\} \\ \frac{\binom{d-3}{j-3}}{\binom{d-1}{j-1}} = \frac{(j-1)(j-2)}{(d-1)(d-2)} & \text{if } k \neq l, i \notin \{k, l\} \\ 0 & \text{if } i \in \{k, l\} \end{cases}.$$

That is, $P(i, j) = \frac{(j-1)(j-2)}{(d-1)(d-2)} (\mathbf{1}_d - e_i) (\mathbf{1}_d - e_i)^T + \frac{(j-1)(d-j)}{(d-1)(d-2)} (I_d - e_i e_i^T)$ and

$$\begin{aligned} \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T &= \frac{j-1}{d-1} (\mathbf{1}_d - e_i) \mathbf{1}_d^T - \frac{(j-1)(j-2)}{(d-1)(d-2)} (\mathbf{1}_d - e_i) (\mathbf{1}_d - e_i)^T - \frac{(j-1)(d-j)}{(d-1)(d-2)} (I_d - e_i e_i^T). \end{aligned}$$

[Step 2] Since $e_{S \cup \{i\}} e_{(S \cup \{i\})^c}^T = (e_S + e_i) (\mathbf{1}_d - e_S - e_i)^T = e_S (\mathbf{1}_d - e_S)^T - e_S e_i^T + e_i (\mathbf{1}_d - e_S)^T - e_i e_i^T$,

$$\begin{aligned} \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_{S \cup \{i\}} e_{(S \cup \{i\})^c}^T &= \frac{j-1}{d-1} (\mathbf{1}_d - e_i) \mathbf{1}_d^T - P(i, j) - \frac{j-1}{d-1} (\mathbf{1}_d - e_i) e_i^T + \frac{d-j}{d-1} e_i (\mathbf{1}_d - e_i)^T \\ &= \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T - \frac{j-1}{d-1} (\mathbf{1}_d - e_i) e_i^T + \frac{d-j}{d-1} e_i (\mathbf{1}_d - e_i)^T. \end{aligned}$$

[Step 3] Therefore, $\Delta_j(x_i) = x^T h_d(i, j) \hat{\beta}$ where

$$\begin{aligned} h_d(i, j) &= \frac{\rho}{1 + \rho(j-1)} \left(\frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T - \frac{j-1}{d-1} (\mathbf{1}_d - e_i) e_i^T + \frac{d-j}{d-1} e_i (\mathbf{1}_d - e_i)^T \right) \\ &\quad + e_i e_i^T - \frac{\rho}{1 - \rho + \rho(j-1)} \left(\frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T \right) \\ &= e_i e_i^T + \left(\frac{\rho}{1 + \rho(j-1)} - \frac{\rho}{1 - \rho + \rho(j-1)} \right) \left(\frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} e_S e_{S^c}^T \right) \\ &\quad + \frac{\rho}{1 + \rho(j-1)} \left(-\frac{j-1}{d-1} (\mathbf{1}_d - e_i) e_i^T + \frac{d-j}{d-1} e_i (\mathbf{1}_d - e_i)^T \right) \\ &= e_i e_i^T + \frac{\rho}{1 + \rho(j-1)} \frac{d-j}{d-1} e_i (\mathbf{1}_d - e_i)^T - \frac{\rho}{1 - \rho + \rho(j-1)} \frac{j-1}{d-1} (\mathbf{1}_d - e_i) e_i^T \end{aligned}$$

$$+ \left(\frac{\rho}{1 + \rho(j-1)} - \frac{\rho}{1 - \rho + \rho(j-1)} \right) \frac{(j-1)(d-j)}{(d-1)(d-2)} \left((\mathbb{1}_d - e_i)(\mathbb{1}_d - e_i)^T - (I_d - e_i e_i^T) \right).$$

□

Now we prove Theorem 1.

Proof of Theorem 1. Without loss of generality, $i \in [d_1]$. That is, a feature of interest is in the first cluster. The key idea of the proof is to only consider the effect of the same cluster. That is, for all possible subsets $S \subseteq [d] \setminus \{i\}$ such that $|S| = j-1$, we can decompose $S = S^{(1)} \cup S^{(2)}$ such that $S^{(1)} \cap S^{(2)} = \emptyset$, $S^{(1)} \cap [d_1] = S^{(1)}$, and $S^{(2)} \cap [d_1] = \emptyset$. In addition, from the Vandermonde's identity, we have

$$\sum_{k=\min\{d_1-j-1-d, 0\}}^{\max\{d_1-1, j-1\}} \binom{d_1-1}{k} \binom{d-d_1}{j-1-k} = \binom{d-1}{j-1}.$$

Hence, the marginal contribution is expressed as

$$\begin{aligned} \Delta_j(x_i) &= \frac{1}{\binom{d-1}{j-1}} \sum_{S \subseteq [d] \setminus \{i\}, |S|=j-1} v^{(\text{cond})}(S \cup \{i\}) - v^{(\text{cond})}(S) \\ &= \frac{1}{\binom{d-1}{j-1}} \sum_{k=\min\{d_1-j-1-d, 0\}}^{\max\{d_1-1, j-1\}} \sum_{S_1 \subseteq [d_1] \setminus \{i\}, |S_1|=k} v^{(\text{cond})}(S_1 \cup \{i\}) - v^{(\text{cond})}(S_1) \\ &= \frac{1}{\binom{d-1}{j-1}} \sum_{k=\min\{d_1-j-1-d, 0\}}^{\max\{d_1-1, j-1\}} \binom{d-d_1}{j-1-k} x_{[d_1]}^T h_{d_1}(i, k+1) \hat{\beta}_{[d_1]} \\ &= x_{[d_1]}^T \left(\frac{1}{\binom{d-1}{j-1}} \sum_{k=\min\{d_1-j-1-d, 0\}}^{\max\{d_1-1, j-1\}} \binom{d-d_1}{j-1-k} h_{d_1}(i, k+1) \right) \hat{\beta}_{[d_1]}. \end{aligned}$$

Here, the second equality is because features in $S^{(2)}$ is independent of $S^{(1)}$ and the last equality is due to Lemma 2. The last term can be expressed as $x^T H(i, j) \hat{\beta}$ by setting a block diagonal matrix

$$H(i, j) = \text{diag} \left(\left(\frac{1}{\binom{d-1}{j-1}} \sum_{k=\min\{d_1-j-1-d, 0\}}^{\max\{d_1-1, j-1\}} \binom{d-d_1}{j-1-k} h_{d_1}(i, k+1) \right), 0_{(d-d_1) \times (d-d_1)} \right).$$

When i -th feature is in m -th cluster, then a similar method gives

$$\begin{aligned} H(i, j) &= \text{diag} \left(0_{(d_1+\dots+d_{m-1}) \times (d_1+\dots+d_{m-1})}, \right. \\ &\quad \left(\frac{1}{\binom{d-1}{j-1}} \sum_{k=\min\{d_m-j-1-d, 0\}}^{\max\{d_m-1, j-1\}} \binom{d-d_m}{j-1-k} h_{d_m}(i, k+1) \right), \\ &\quad \left. 0_{(d-d_1-\dots-d_m) \times (d-d_1-\dots-d_m)} \right). \end{aligned}$$

It concludes a proof. □

B The Shapley value and the semivalue

In this section, we elaborate on the Shapley value and the semivalue with a focus on their mathematical properties. To this end, we denote an attribution that is based on a coalition function v by $\phi(\cdot; v)$ ¹.

Shapley axioms We first introduce the four Shapley axioms: Linearity, Null player, Symmetry, and Efficiency.

- **Linearity:** for coalition functions $v_1, v_2 \in \{u \mid u : 2^{[d]} \rightarrow \mathbb{R}\}$ and $\alpha_1, \alpha_2 \in \mathbb{R}$, $\phi(x_i, \alpha_1 v_1 + \alpha_2 v_2) = \alpha_1 \phi(x_i, v_1) + \alpha_2 \phi(x_i, v_2)$.
- **Null player:** if $v(S \cup \{i\}) = v(S) + c$ for every $S \subseteq [d] \setminus \{i\}$ and some $c \in \mathbb{R}$ and $v : 2^{[d]} \rightarrow \mathbb{R}$, then $\phi(x_i; v) = c$.
- **Symmetry:** for every $v : 2^{[d]} \rightarrow \mathbb{R}$ and every permutation π on $[d]$, $\phi(x_i; \pi^* v) = \pi^* \phi(x_i; v)$ where $\pi^* v$ is defined as $(\pi^* v)(S) := v(\pi(S))$ for every $S \subseteq [d]$.
- **Efficiency:** for every $v : 2^{[d]} \rightarrow \mathbb{R}$, $\sum_{i \in [d]} \phi(x_i; v) = v([d])$.

Shapley [1953a] showed that the Shapley value is the unique function that satisfies the four axioms, and it has been studied in the literature [Owen, 2014, Lundberg and Lee, 2017].

Semivalue The relaxation of the Shapley axioms has been a central topic in cooperative game theory [Monderer and Samet, 2002, Weber, 1988]. Our WeightedSHAP is based on the semivalue [Dubey and Weber, 1977], which relaxes the Efficiency axiom. We formally define the semivalue in the following definition.

Definition 3 (semivalue). *We say a function ϕ is a semivalue if ϕ satisfies Linearity, Null player, and Symmetry axioms.*

In the machine learning literature, Kwon and Zou [2021] used the semivalue to quantify the importance of individual data points. In addition, they showed the theoretical properties of the semivalue in the language of the data valuation problem. In the following, we provide the counterpart theorems on the attribution problem settings.

Proposition 3 (Theorem 2 of Kwon and Zou [2021]). *An attribution ϕ is a semivalue, if and only if, there exists a weight vector $\mathbf{w} = (w_1, \dots, w_d)^T$ such that $\sum_{j=1}^d \binom{d-1}{j-1} w_j = d$ and the attribution ϕ can be expressed as follows.*

$$\phi(x_i; v, \mathbf{w}) := \frac{1}{d} \sum_{j=1}^d \binom{d-1}{j-1} w_j \Delta_j(x_i; v). \quad (7)$$

By setting $\tilde{w}_j = \binom{d-1}{j-1} w_j / d$, it can be represented as a weighted mean of the marginal contribution in Equation 5, i.e., $\phi_{\tilde{\mathbf{w}}} = \phi(x_i; v, \mathbf{w})$.

Proposition 4 (Proposition 3 of Kwon and Zou [2021]). *Let ϕ_1 and ϕ_2 be two attribution methods such that for any v , the sum of feature attributions are same, i.e.,*

$$\sum_{i=1}^d \phi_1(x_i; v) = \sum_{i=1}^d \phi_2(x_i; v).$$

Then, the two attribution methods are identical, i.e., $\phi_1 = \phi_2$.

That is, if there are two attributions with the same total sum of attributions across all coalition function v , then they are identical. Although the semivalue is not unique, but it is *almost* unique in that it is the only semivalue with a particular sum of attributions.

We deploy the concept of the semivalue into the attribution problem, which makes the main difference from Kwon and Zou [2021]. In contrast to the data valuation problem, where the marginal contribution based on a small coalition size is preferred because it is more effective to capture the label errors, we observe that it is not necessarily true in the attribution problem. As we discussed in the manuscript, the marginal contribution based on a large coalition size Δ_d can be preferred due to the signal, but a weighted mean can be preferred to reduce an estimation error.

¹In the manuscript we omitted $v^{(\text{cond})}$ for notational convenience.

C Implementation details

In this section, we provide implementation details used in Example in Section 4 and numerical experiments in Section 5. We present implementation algorithms in Section C.1, datasets in Section C.2, and experimental settings in Section C.3. Our Python-based implementations are available at <https://github.com/ykwon0407/WeightedSHAP>.

C.1 Implementation algorithm

Given a finite set \mathcal{W} , an easy-to-compute utility function \mathcal{T} and the marginal contribution estimates $\Delta_j(x_i)$, the optimal weight \mathbf{w}^* can be achieved by iteratively evaluating the utility \mathcal{T} for each attribution method $\phi_{\mathbf{w}}$ with $\mathbf{w} \in \mathcal{W}$. This procedure is described in Algorithm 1.

Algorithm 1 Computation of WeightedSHAP

Require: A finite set of weights \mathcal{W} . A utility function \mathcal{T} . A set of estimates for the marginal contributions $\Delta_j(x_i)$ obtained from Algorithm 2.

procedure

Initialize a constant $C_{\mathcal{T}} = -\infty$.

for $\mathbf{w} \in \mathcal{W}$ **do**

Compute $\phi_{\mathbf{w}}(x_i) = \sum_{j=1}^d \mathbf{w}_j \Delta_j(x_i)$ for all $i \in [d]$.

Rank features $(x_{(1)}, \dots, x_{(d)})$ based on the absolute value of their attribution, *i.e.*, $|\phi_{\mathbf{w}}(x_{(1)})| \geq \dots \geq |\phi_{\mathbf{w}}(x_{(d)})|$.

Evaluate $\mathcal{T}(\phi_{\mathbf{w}}; x, \hat{f})$.

if $C_{\mathcal{T}} \leq \mathcal{T}(\phi_{\mathbf{w}}; x, \hat{f})$ **then**

$C_{\mathcal{T}} \leftarrow \mathcal{T}(\phi_{\mathbf{w}}; x, \hat{f})$ and $\mathbf{w}^*(\mathcal{T}, \mathcal{W}) \leftarrow \mathbf{w}$.

end if

end for

$\phi_{\text{WeightedSHAP}}(\mathcal{T}, \mathcal{W}) := \phi_{\mathbf{w}^*(\mathcal{T}, \mathcal{W})}$

end procedure

Algorithm 2 Estimation of the marginal contributions $\Delta_j(x_i)$

Require: A conditional coalition estimate $\hat{v}^{(\text{cond})}$. A terminating threshold ρ .

procedure

Initialize $\hat{\rho} = 2\rho$, $B = 1$, $\Delta_j^{(0)}(x_i) = 0$ for all $i, j \in [d]$.

while $\hat{\rho} \geq \rho$ **do**

for $i \in [d]$ **do**

$S \leftarrow \emptyset$.

Generate a random order of $[d] \setminus \{i\}$ and denote it by η .

for $j \in [d]$ **do**

$\Delta_j^{(B)}(x_i) \leftarrow \frac{B-1}{B} \Delta_j^{(B-1)}(x_i) + \frac{1}{B} (\hat{v}^{(\text{cond})}(S \cup \{i\}) - \hat{v}^{(\text{cond})}(S))$.

$S \leftarrow S \cup \eta(j)$.

end for

end for

Compute the Gelman-Rubin statistics for $\{\Delta_j^{(b)}(x_i)\}_{b=1}^B$ and take its maximum value among $i, j \in [d]$.

$B \leftarrow B + 1$.

end while

end procedure

Throughout our experiments, \mathcal{T} is the negative value of the AUP defined in 4 and \mathcal{W} is a set of 13 different weights, namely, $\mathcal{W} = \{\Delta_1, \Delta_d\} \cup \{\phi_{\mathbf{w}^{\text{Beta}}(\alpha, \beta)} \mid (\alpha, \beta) \in \{(16, 1), (8, 1), (4, 1), (2, 1), (1, 1), (1, 2), (1, 4), (1, 8), (1, 16), (1, 32)\}\}$. Here, $\mathbf{w}^{\text{Beta}}(\alpha, \beta) \in [0, 1]^d$ is defined as

$$(\mathbf{w}^{\text{Beta}}(\alpha, \beta))_j := \binom{d-1}{j-1} \frac{\text{Beta}(j + \beta - 1, d - j + \alpha)}{\text{Beta}(\alpha, \beta)},$$

where $\text{Beta}(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha, \beta)$ is the Beta function and $\Gamma(\cdot)$ is the Gamma function. Inspired by Kwon and Zou [2021], the key motivation of this function is that (i) this type of a functional form is known to satisfy the condition in Proposition 3 and (ii) it has a closed-form expression, so it is easy-to-compute. Note that $\mathbf{w}^{\text{Beta}}(1, 1) = d^{-1}\mathbb{1}_d$ generates the Shapley value, and our algorithm guarantees the superior performance of WeightedSHAP to the Shapley value as $\mathbf{w}^{\text{Beta}}(1, 1) \in \mathcal{W}$. Since the Beta weight vector $\mathbf{w}^{\text{Beta}}(\alpha, \beta)$ distributes weights throughout all the marginal contributions $\Delta_1, \dots, \Delta_d$, we can think of it as a regularized version of $\mathbb{1}(j = 1)$ and $\mathbb{1}(j = d)$, which generates Δ_1 and Δ_d , respectively.

The key input of Algorithm 1 is the marginal contribution estimates. To obtain this estimate, we first need to estimate a conditional coalition function $v^{(\text{cond})}$. Following the literature [Frye et al., 2020, Jethani et al., 2021a,b], we obtain this by training a surrogate model² that takes as input a subset of input features and outputs a conditional expectation of a prediction value given the same subset. As for the surrogate model, we use a multilayer perceptron model with the two hidden layers, and each layer has 128 neurons with the ELU nonlinear activation function [Clevert et al., 2015]. For the classification (*resp.* regression) tasks, we consider the Kullback-Libeler divergence (*resp.* MSE loss) as a loss function as suggested in [Jethani et al., 2021a]. We use the held-out dataset to learn a surrogate model, which corresponds to 10% of the original dataset. Throughout our experiments, we mostly follow the hyperparameters used in the previous work Jethani et al. [2021a]. For instance, we use the same Adam optimizer [Kingma and Ba, 2014] with the initial learning rate of 10^{-3} , the epochs of 100, and the mini-batch size of 64.

With the estimator $\hat{v}^{(\text{cond})}$, the marginal contributions are estimated as in Algorithm 2. Using the fact that the marginal contribution in Equation (2) is defined as a simple average over all the possible subsets S with the same coalition size, we use a sampling-based algorithm [Ghorbani and Zou, 2019, Kwon and Zou, 2021]. It is well known that a sampling-based algorithm guarantees the convergence to the true marginal contribution value when sampling procedures are repeated. In our experiments, we use a finite number of samplings based on the Gelman-Rubin stopping criteria [Vats and Knudson, 2021, Equation (4)]. We regard the repeated sampling procedures as 10 Markov chains and compute the Gelman-Rubin statistic for the marginal contribution $\Delta_j(x_i)$ for all $i, j \in [d]$. We take their maximum for every iteration and stop the sampling procedure if this maximum is smaller than a prefixed terminating threshold. We use a terminating threshold $\rho = 1.005$ which is much smaller than a typical terminating threshold 1.1 [Gelman et al., 1995].

C.2 Datasets

In Section 5 and D.3, the four real-world datasets are used for the regression tasks. The one synthetic dataset (`gaussian`) and the seven real-world datasets are used for the classification tasks. All the real-world datasets are downloaded from the UCI repository [Dua and Graff, 2017] or the OpenML platform³. As for the synthetic classification dataset (`gaussian`), we consider the following distribution. For $d = 30$ and $\rho = 0.25$,

$$X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{(\rho, d)}), \quad Y = \text{Bern}(p(X)),$$

where $p(X) := \exp(X^T \beta^*) / (1 + \exp(X^T \beta^*))$ for $\beta^* = (1, 0.98, \dots, 0.82, 0, \dots, 0)^T$. The first 10 features are associated with an output and the last 20 features are not directly associated.

As for the real datasets, from a given raw dataset, we only use the continuous features and exclude an observation if some information is missing. Table 1 summarizes the basic information after preprocessing the raw datasets. From this dataset, we randomly take 10,000 samples and consider it as the entire dataset. We normalize every feature to have zero mean and unit variance. In addition, we generate spurious features that are associated with the original features, but not directly associated with an output by its construction. We repeat the following procedure until the number of features becomes three times the original number of features: given an input matrix $X \in \mathbb{R}^{n \times p}$, we generate a new column $X_{\text{new}} \in \mathbb{R}^{n \times 1}$ by $X_{\text{new}} = \frac{\rho}{1 + \rho(p-1)} X \mathbb{1}_p + \sqrt{1 - \frac{\rho^2 p}{1 + \rho(p-1)}} \varepsilon$ where $\varepsilon \sim \mathcal{N}(0, 1)$ is randomly drawn from a standard Gaussian distribution and ρ is the average of off-diagonal terms in the correlation matrix of the original input matrix. This procedure is to append more features while

²It is known that this surrogate model unbiasedly estimates a conditional expectation of a prediction value given a subset of features under mild conditions [Frye et al., 2020, Covert et al., 2020].

³Website: <https://www.openml.org/>

Table 1: A summary of real-world datasets used in experiments.

Dataset	Sample size	Input dimension	Source
Regression			
abalone	4177	10	UCI Repository
boston	506	13	Harrison Jr and Rubinfeld [1978]
airfoil	1503	5	UCI Repository
whitewine	4898	11	UCI Repository
Classification			
cpu	8192	22	https://www.openml.org/d/761
fraud	284807	31	Dal Pozzolo et al. [2015]
phoneme	5404	6	https://www.openml.org/d/1489
wind	6574	15	https://www.openml.org/d/847
adult	32561	11	UCI Repository
2dplanes	40768	11	https://www.openml.org/d/727
click	1997410	12	https://www.openml.org/d/1218

keeping the average correlation between two features same. We observe several preprocessing steps we deploy do not affect our main result that WeightedSHAP consistently assigns large values for more influential features. However, they are deployed to increase the stability in training a prediction model and to reflect practical data analysis situations (e.g., features are highly correlated, but not necessarily every feature is associated with an output).

C.3 Experimental settings

Experiments in Section 3.2 As we have closed-form expressions under the Gaussian assumption by Theorem 1, we do not use any estimation algorithm here. As for the Figure 1(a), explicit forms for the Shapley value and the optimal order have been used. As for the Figure 1(b), we first generate 10,000 training data points as follows. For $d = 100$ and $\rho \in \{0.2, 0.6\}$,

$$X \sim \mathcal{N}(\mathbf{0}_d, \Sigma_{(\rho, d)}), \quad Y = X^T \beta^* + 2 \times \varepsilon,$$

where $\varepsilon \sim \mathcal{N}(0, 1)$ is a random Gaussian error and $\beta^* = (1, 0.99, \dots, 0.81, 0, \dots, 0)^T$. That is, the first 20 features are associated with an output, but the last 80 features are not directly associated with the output. We train a linear model based on the 10,000 samples and evaluate feature attributions for the held-out 100 samples. All the results in Figure 1(b) is based on this held-out test dataset.

Experiment in Section 4 In contrast to the experiment in Section 3.2, we estimate the marginal contributions following Algorithm 2. Other experiment settings are same as in Section 3.2.

Experiments in Section 5 We first split the original dataset into four parts: a training dataset, a validation dataset, a dataset to obtain a surrogate model, and a test dataset. The ratio between them is 70%, 10%, 10%, and 10%, but we take $n_{\text{test}} := \max(0.1 \times N, 100)$ samples for a test dataset. Here N denotes the number of the original sample size. The training and validation datasets are used for a prediction model \hat{f} , and a dataset to obtain a surrogate model is used for $\hat{v}^{(\text{cond})}$. After that we compute WeightedSHAP for n_{test} samples. All the results (e.g., prediction recovery error, MSE, and AUC) are based on these n_{test} test samples.

For a boosting model, we use the `lightgbm` algorithm [Ke et al., 2017] with the learning rate of 0.005 and 15 final leaves and apply the early stopping with the 25 patience epoch. For the classification (*resp.* regression) tasks, we use the cross entropy (*resp.* MSE) loss function. As for the estimation of WeightedSHAP, see Section C.1. In Section D.2, we simply use the least squares estimator for a linear prediction model.

D Additional experimental results

In this section, we provide additional experimental results. Our main findings presented in the manuscript are consistently observed on different models and datasets. We first present estimation error analysis on a different dataset in Section D.1, additional experimental results when a prediction model is linear in Section D.2, and additional experimental results on different classification datasets in Section D.3.

D.1 Estimation error analysis when $\rho = 0.2$

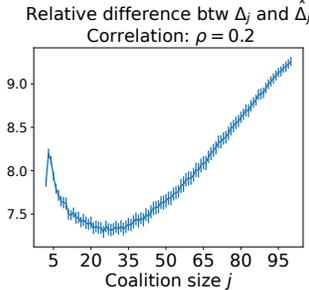


Figure 6: Illustrations of the relative difference between the true marginal contribution Δ_j and its estimate $\hat{\Delta}_j$ as a function of the coalition size $j \in [100]$. We consider the same setting used in Figure 1(b) with $\rho = 0.2$. The Δ_{100} is shown to have the largest relative difference.

We provide an additional estimation error result with a different dataset. In contrast to the Figure 2 shown in the manuscript, we generate data points from the Gaussian distribution $\mathcal{N}(\mathbf{0}_{100}, \Sigma_{(0.2, 100)})$, *i.e.*, $\rho = 0.2$.

Figure 6 shows the relative difference between the true marginal contribution Δ_j and its estimate $\hat{\Delta}_j$ as a function of the coalition size $j \in \{1, \dots, 100\}$. Similar to Figure 2, the relative difference of Δ_{100} has the largest estimation error. Given that Δ_{100} is the most informative marginal contribution in Figure 1(b), it again suggests the use of a weighted mean of the marginal contributions.

D.2 Additional experimental results when a prediction model is linear

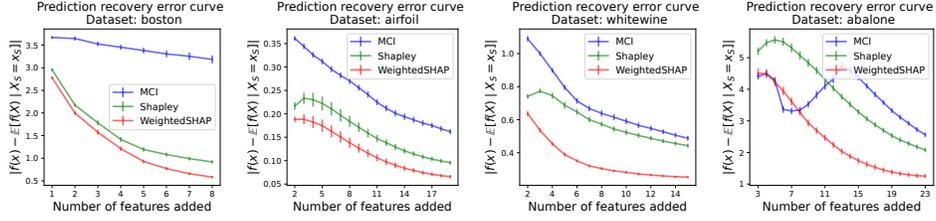
We conduct additional analyses when a prediction model $\hat{f}(x)$ is linear. We use the same experimental setting as in Section 5, but only a prediction model is changed from a boosting model to a linear model.

Figures 7(a) and 7(b) (*resp.* Figures 7(c) and 7(d)) compare WeightedSHAP with the Shapley on the regression tasks (*resp.* the classification tasks). In all experiments, the prediction recovery error of WeightedSHAP is significantly smaller than the MCI and the Shapley value. Moreover, the performance metrics, namely MSE and AUC, of WeightedSHAP are better than or at least comparable to the MCI and the Shapely value. In short, our findings are consistently observed with a linear prediction model, showing the influential features identified by WeightedSHAP are better able to recapitulate the model’s predictions compared to the features identified by baseline attribution methods.

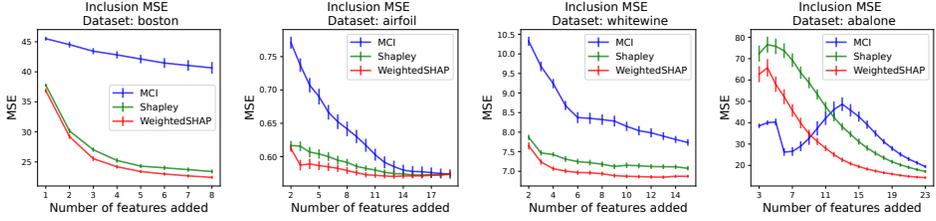
D.3 Additional experimental results with different classification datasets

We conduct additional experiments on different classification datasets. A boosting model is used for a prediction model. In terms of the experimental setting, the only difference from Figures 4(a) and 4(b) is the datasets. Details on the four datasets are provided in Section C.2.

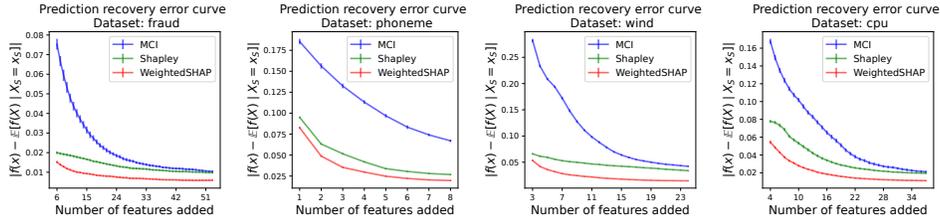
Figures 8(a) and 8(b) show additional experimental results on the four different classification datasets. As in the previous experiments, the prediction recovery error of WeightedSHAP is lower than the Shapley value. As for the inclusion AUC metric, WeightedSHAP is significantly better or at least comparable to the Shapley value.



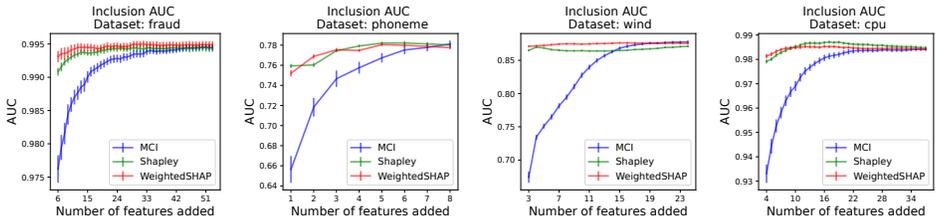
(a) Illustrations of the prediction recovery error curve on the four regression datasets. The lower, the better.



(b) Illustrations of the Inclusion MSE curve on the four regression datasets. The lower, the better.



(c) Illustrations of the prediction recovery error curve on the four binary classification datasets. The lower, the better.

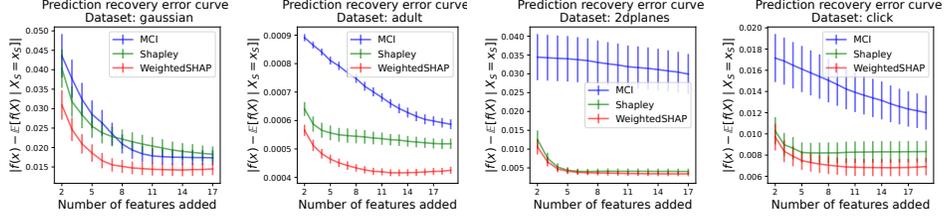


(d) Illustrations of the Inclusion AUC curve on the four binary classification datasets. The higher, the better.

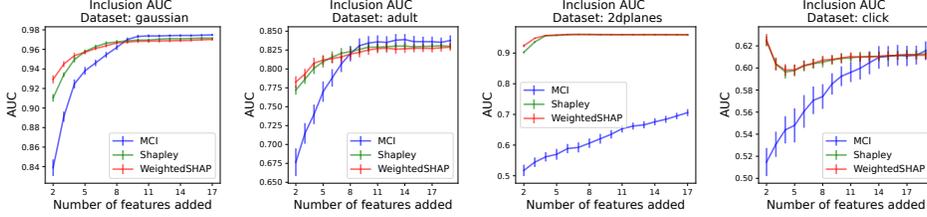
Figure 7: **When a linear prediction model is used.** Illustrations of the prediction recovery error curve and the performance curve as a function of the number of features added. We add features from most influential to the least influential. We denote a 95% confidence interval based on 30 independent runs. Our main findings are consistently observed with a linear prediction model.

D.4 Additional experimental results with different evaluation metrics

We further examine the performance of WeightedSHAP on two different evaluation metrics: the Exclusion performance used in Jethani et al. [2021a] and the Inclusion performance using masked features used in Masoomi et al. [2021]. As for the Inclusion performance in Figures 9(b) and 9(d), the main difference from Figures 3(b) and 4(b) is that it considers model predictions based on mean-masked features whenever new features are added, not the conditional expectation. For each evaluation metric, we use different utility functions in the weight optimization (6). Specifically, for



(a) Illustrations of the prediction recovery error curve on the four binary classification datasets. The lower, the better.



(b) Illustrations of the Inclusion AUC curve on the four binary classification datasets. The higher, the better.

Figure 8: **Additional results on different classification datasets.** Illustrations of the prediction recovery error curve and the performance curve as a function of the number of features added. We add features from most influential to the least influential. We denote a 95% confidence interval based on 50 independent runs. Our main findings are consistently observed on various classification datasets.

the Exclusion task, we use

$$\mathcal{T}(\phi) = \sum_{k=1}^d \left| \hat{f}(x) - \mathbb{E}[\hat{f}(X) \mid X_{\mathcal{J}(k;\phi,x)} = x_{\mathcal{J}(k;\phi,x)}] \right|,$$

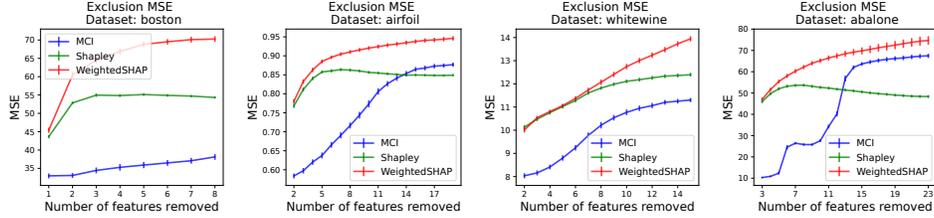
where $\mathcal{J}(k; \phi, x) := [d] \setminus \mathcal{I}(k; \phi, x)$ and $\mathcal{I}(k; \phi, x) \subseteq [d]$ be a set of k integers that indicates k most influential features based on their absolute value $|\phi(x_j)|$. That is, it adds the least important feature first. This is equivalent to removing the most important feature from the entire set of features. As for the Inclusion task, we consider the following quantity.

$$\mathcal{T}(\phi) = \sum_{k=1}^d \left| \hat{f}(x) - \hat{f}(x_{\mathcal{I}(k;\phi,x)}, \mu_{\mathcal{J}(k;\phi,x)}) \right|,$$

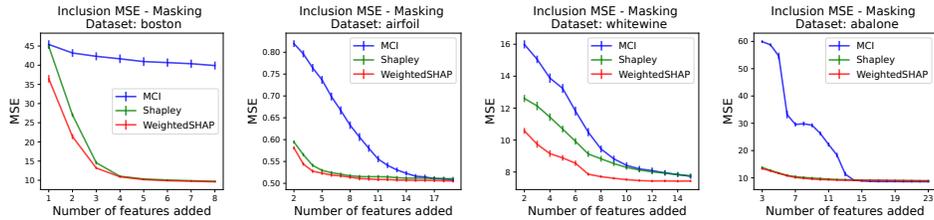
where μ is a pre-computed mean of the entire features. The main idea is that μ part is non-informative to a particular prediction, but after replacing the μ part with influential features, the quantity $\left| \hat{f}(x) - \hat{f}(x_{\mathcal{I}(k;\phi,x)}, \mu_{\mathcal{J}(k;\phi,x)}) \right|$ is expected to be reduced.

We note that there is not an agreed-upon objective metric for ML interpretability, and at the same time, practitioners always can choose their downstream objectives depending on their tasks. Given that the concept of influential features can be dependent on downstream tasks, we believe the attribution should be optimized to downstream objectives. For instance, there is no single attribution method that works universally well simultaneously on the Inclusion AUC and Exclusion AUC tasks, and the optimal attribution depends on the evaluation metric. Our WeightedSHAP is anticipated to be flexible in optimizing downstream objectives and finds the most suitable attributions. In contrast, the Shapley value is fixed to every downstream task.

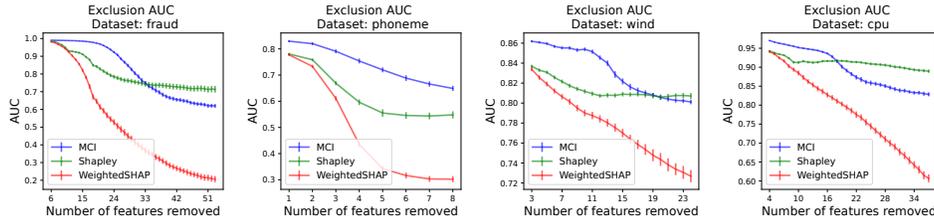
Figures 9(a) and 9(b) (*resp.* Figures 9(c) and 9(d)) compare MSE (*resp.* AUC) curves of the three different attribution methods. In most experimental settings, as anticipated, WeightedSHAP substantially outperforms the MCI and the Shapley value. With flexible choices of the utility function, WeightedSHAP shows significantly better performances than standard attribution methods across different downstream objectives.



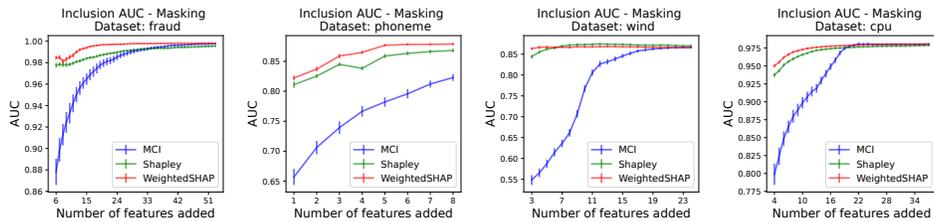
(a) Illustrations of the Exclusion MSE on the four regression datasets. The higher, the better.



(b) Illustrations of the Inclusion MSE curve using masked features on the four regression datasets. The lower, the better.



(c) Illustrations of the Exclusion AUC curve on the four binary classification datasets. The lower, the better.



(d) Illustrations of the Inclusion AUC curve using masked features on the four binary classification datasets. The higher, the better.

Figure 9: **WeightedSHAP on different evaluation metrics.** Illustrations of the Exclusion performance curve (*resp.* Inclusion performance curve using masked features) as a function of the number of removed (*resp.* added) features from most influential to the least influential. We denote a 95% confidence interval based on 50 independent runs. WeightedSHAP significantly outperforms the MCI and the Shapley value on different evaluation metrics.