

Shaping Up SHAP: Enhancing Stability through Layer-Wise Neighbor Selection

Gwladys Kelodjou¹, Laurence Rozé², Véronique Masson¹, Luis Galárraga¹,
Romaric Gaudel¹, Maurice Tchuente³, Alexandre Termier¹

¹Univ Rennes, Inria, CNRS, IRISA - UMR 6074, F35000 Rennes, France

²Univ Rennes, INSA Rennes, CNRS, Inria, IRISA - UMR 6074, F35000 Rennes, France

³Sorbonne University, IRD, University of Yaoundé I, UMI 209 UMMISCO, P.O. Box 337 Yaoundé, Cameroon
gwladys.kelodjou@irisa.fr, laurence.roze@irisa.fr

Abstract

Machine learning techniques, such as deep learning and ensemble methods, are widely used in various domains due to their ability to handle complex real-world tasks. However, their black-box nature has raised multiple concerns about the fairness, trustworthiness, and transparency of computer-assisted decision-making. This has led to the emergence of local post-hoc explainability methods, which offer explanations for individual decisions made by black-box algorithms. Among these methods, Kernel SHAP is widely used due to its model-agnostic nature and its well-founded theoretical framework. Despite these strengths, Kernel SHAP suffers from high instability: different executions of the method with the same inputs can lead to significantly different explanations, which diminishes the relevance of the explanations. The contribution of this paper is two-fold. On the one hand, we show that Kernel SHAP's instability is caused by its stochastic neighbor selection procedure, which we adapt to achieve full stability without compromising explanation fidelity. On the other hand, we show that by restricting the neighbors generation to perturbations of size 1 – which we call the coalitions of Layer 1 – we obtain a novel feature-attribution method that is fully stable, computationally efficient, and still meaningful.

Introduction and Related Work

Modern machine learning techniques, such as deep learning and ensemble methods, have become invaluable tools in diverse high-stakes fields such as healthcare and banking, due to their accuracy at handling complex real-world tasks. Yet, their “black-box” nature has raised concerns about their fairness and trustworthiness, which in turns has motivated the interest in explainability. In particular, *local post-hoc* explainability is concerned with providing a posteriori explanations for the decisions of black-box algorithms on single instances. Notable approaches in this line are LIME (Ribeiro, Singh, and Guestrin 2016), SHAP (Lundberg and Lee 2017), and Grad-CAM (Selvaraju et al. 2017), which compute feature-attribution explanations. A feature attribution explanation assigns a contribution score to each of the model's input features. This tells users how (positively or negatively) and how much an input feature influenced the black-box's answer on a target instance.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Among the explainability methods based on feature attribution, SHAP is particularly popular because its attribution scores, the **SHAP values**, are inspired by the well-founded **Shapley values** (Shapley et al. 1953) from coalitional game theory. In practice, SHAP implementations typically compute approximations of the theoretical SHAP values because of the exponential complexity of the exact calculation. For instance, the widely used **Kernel SHAP** – SHAP's model-agnostic flavor – estimates the attribution scores by solving a linear regression on a random sample of *neighbor instances*, constructed by perturbing the target instance.

While local post-hoc explanations based on perturbations are widely used to explain black-box models (LIME also falls in this category), this class of methods are known to suffer from instability. Zhou, Hooker, and Wang (2021) argue that a key property of any explanation technique is its *stability* or reproducibility, defined as the method's ability to deliver the same results under the same conditions across multiple executions. They therefore propose a framework to determine the number of perturbation points required to ensure stable LIME explanations. Visani et al. (2022) introduced two metrics to quantify the stability of LIME explanations: the Variable Stability Index, which measures the variation of the reported explanatory features, and the Coefficient Stability Index, which captures the variation of the coefficients associated to those features. Visani, Bagli, and Chesani (2020) offer a framework that allows users to choose the level of stability and fidelity for LIME explanations.

We highlight the difference between *stability* and the criterion called *robustness* introduced by Alvarez-Melis and Jaakkola (2018). An explanation method is robust if similar, but not identical, inputs entail similar explanations. Alvarez-Melis and Jaakkola (2018) quantify robustness as the maximum change in explanations obtained when small perturbations are applied to a given instance. They demonstrate that state-of-the-art methods (including LIME and SHAP) do not satisfy the robustness criterion. Stability and robustness are different concepts, however stability is a necessary condition for robustness.

While stability has been well studied for LIME, and stable improvements of LIME have been proposed, this is not the case for SHAP. In this paper, we show that the Kernel SHAP's estimator also suffers from stability issues, which stem from its stochastic neighbor selection. Given the popu-

larity of Kernel SHAP, such instability diminishes the utility of SHAP, not to mention its potential impact on the trust in post-hoc explainability approaches. We list the contributions of this paper in the following.

- Our first contribution is an alternative procedure for choosing the neighbors of a target instance in Kernel SHAP so that stability is improved.
- Our second contribution is a thorough experimental study of the stability but also the fidelity of Kernel SHAP under different neighbor selection procedures. It shows that our proposed neighbor selection approach achieves stability without hurting explanation fidelity.
- Motivated by the results of our experiments, our third contribution is a novel approximation of the SHAP values, which relies on neighbors generated from simple perturbations – called the *layer-1 neighbors*. We demonstrate, via a theoretical study, that the obtained scores meet the main properties of attribution methods based on coalitional game theory. This means that our approach is an attribution method in itself, which provides scores that are stable and faster to compute than the SHAP values, while retaining their good properties.

The remainder of this paper is organized as follows: in the *Background* section, we provide an overview of SHAP and Kernel SHAP. The *Reason for Instability* section brings to light the instability issues of Kernel SHAP. In the *Improving Kernel SHAP’s Stability* section, we present our approach for improving the stability of Kernel SHAP. The *Experiments* section provides an empirical evaluation of our proposed method. This experiment shows that using only *layer-1* neighbors already yields good results. Based on these results, in the *First Layer Attribution Analysis* section, we conduct a theoretical analysis on the properties of our attribution method based on *layer-1* neighbors. We evaluate our attribution method in the latter part of this section by comparing the attributions of the exact SHAP values with those of *layer-1*. Finally, we conclude the paper in the last section, summarizing our findings and suggesting future research directions.

Background

We now review the definitions of Shapley and SHAP values.

Shapley Values

The Shapley value is a mathematical concept from game theory that was introduced by Shapley et al. (1953) to allocate performance credit across coalition game players. It is used to evaluate the *fair share* of a player in a cooperative game (Ghorbani and Zou 2019). Suppose that we have a cooperative game with M players numbered from 1 to M , and let N denote the set of players: $N = \{1, 2, \dots, M\}$. A coalitional game is a tuple $\langle N, v \rangle$ where $v : 2^N \rightarrow \mathbb{R}$ is a characteristic function such that $v(\emptyset) = 0$ (Strumbej and Kononenko 2010). Subsets of N are **coalitions** and v returns a value for each coalition S ($S \subseteq N$). So, for each coalition S , $v(S)$ represents the reward that arises when the players in the subset S participate in the game. $v(S)$ can be seen as the worth of coalition S . The set N is also a coalition, and we call it

the **grand coalition**. Thus $v(N)$ is the total payoff that we aim to split among the players in a *fair way*. Fair allocations must be based on each player’s contribution to the profit. The marginal contribution of a player j to a coalition S is defined as the additional value induced by including j in the coalition: $v(S \cup \{j\}) - v(S)$. As we can see, this marginal contribution depends on the composition of the coalition. To determine the player’s overall contribution, it is necessary to consider all possible subsets of players, i.e., coalitions. Then, the Shapley value $\phi_j(v)$ of player j for a game v is calculated as a weighted average of the marginal contributions to the coalitions in which that player participates (Ruiz, Valenciano, and Zarzuelo 1998):

$$\phi_j(v) = \sum_{S \subseteq N \setminus \{j\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} [v(S \cup \{j\}) - v(S)]. \quad (1)$$

Shapley et al. (1953) proved that this value is the unique allocation of the grand coalition which satisfies the following axioms: Efficiency, Symmetry, Dummy, and Additivity.

In recent years, the concept of Shapley value has become a popular technique for interpreting machine learning models (Song, Nelson, and Staum 2016; Ghorbani and Zou 2019; Lundberg and Lee 2017). In this setting, the players are the features used by the model and the gain is the model’s prediction. The aim is therefore to assign each feature an importance value that represents its effect on the model’s prediction. Let f be the predictive model to be explained, and $N = \{1, 2, \dots, M\}$ be a set representing M input features. f maps an input feature vector $x = [x_1, x_2, \dots, x_M]$ to an output $f(x)$ where $f(x) \in \mathbb{R}$ in regression and $f(x) \in [0, 1]^{|C|}$ in classification, with C being a finite set of labels. Shapley’s equation thus becomes:

$$\phi_j(f, x) = \sum_{S \subseteq N \setminus \{j\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} [f_{S \cup \{j\}}(x_{S \cup \{j\}}) - f_S(x_S)], \quad (2)$$

where x_S represents the sub-vector of x restricted to the feature subset S , and $f_S(x_S)$ is the model output restricted to the feature subset S . To compute the Shapley values in such a setting, we must find out how to calculate $f_S(x_S)$, because most models cannot handle arbitrary patterns of missing input values – they are only defined for the grand coalition. There are several approaches to do so. Here we will look at how SHAP (Lundberg and Lee 2017) works.

SHAP Values

SHAP values (Lundberg and Lee 2017) are defined as the coefficients of an additive surrogate explanation model which is a linear function of binary features:

$$g(z') = \phi_0 + \sum_{i=1}^M \phi_i z'_i. \quad (3)$$

Here, $z' \in \{0, 1\}^M$ is a binary vector representing a coalition of features, where $z'_i = 1$ when the i -th feature is observed and $z'_i = 0$ otherwise. The values $\phi_i \in \mathbb{R}$ correspond to the feature attribution values.

For example, a coalition $z' = [1, 0, 0, 1]$ simply represents the coalition $S = \{x_1, x_4\}$ if we assume $M = 4$. Equation (4) gives the Shapley's values, where $|z'|$ is the number of 1s in vector z' , and $z' \setminus i$ is derived from z' by setting $z'_i = 0$.

$$\phi_i(f, x) = \sum_{z' \in \{0,1\}^M} \frac{(|z'| - 1)!(M - |z'|)!}{M!} [f_x(z') - f_x(z' \setminus i)] \quad (4)$$

In Equation (4), $f_x(z') = \mathbb{E}[f(z) | z_S = x_S]$ with $S = \{i \in N | z'_i = 1\}$ and z_S the sub-vector of x restricted to the feature subset S . In other words $f_x(z')$ is the average of f 's answers when we fix the present features in z' , and consider all possible values of the absent features. SHAP values satisfy three desirable properties: *local accuracy*, *missingness* and *consistency* (Lundberg and Lee 2017), which are equivalent to those of the Shapley values (regarding local accuracy and consistency). The exact computation of the SHAP values is exponential in the number of features, so numerous approaches have been proposed to approximate them. We are interested here in **Kernel SHAP**, which is the main model-agnostic variant of SHAP.

Kernel SHAP. It is a linear-regression-based approximation of the SHAP values. The idea is to find a surrogate function g by solving a weighted least squares problem with weighting kernel π_x :

$$\arg \min_g \sum_{z' \in \{0,1\}^M} \pi_x(z')(g(z') - f_x(z'))^2 \quad (5)$$

$$\pi_x(z') = \frac{M - 1}{\binom{M}{|z'|} |z'| (M - |z'|)} \quad (6)$$

where g in Equation (5) is defined as in Equation (3). We note that the values $\pi_x(z') = \infty$ when $|z'| \in \{0, M\}$. This enforces constraints (i) $\phi_0 = f_x(\emptyset) = E[f(x)]$ for the intercept of g and (ii) $\sum_{i=1}^M \phi_i = f(x) - f_x(\emptyset)$ for the sum of the coefficients. The latter constraint is the *local accuracy* constraint, which states that the sum of the attribution coefficients equals the total payoff, the prediction $f(x)$ in this case. To approximate the SHAP values, Kernel SHAP samples a subset of coalitions and then uses the sample to solve Equation (5). The size of the sample used to learn Kernel SHAP is called the **budget**. Given a budget, the sampling is guided by the coalition weights (π_x , also called *Shapley Kernel*). The weight assigned to each coalition is determined by the number of features involved in that coalition (denoted as $|z'|$ and representing the number of 1s in z'). From now we use the terms neighbor and coalition interchangeably.

Layer i is defined as the set of coalitions having i features present or i features absent (or $M - i$ features present). All the coalitions within the same layer have the same weight. Table 1 is an example of some coalitions and their corresponding layers for $M = 4$. As i approaches $\lfloor \frac{M}{2} \rfloor$, the weight assigned to the coalitions within that layer decreases. Consequently, the sampling procedure progresses incrementally, starting from lower-level layers and gradually moving towards higher-level layers. Specifically, this involves generating first the coalitions with 1 or $M - 1$ features included. If the budget is not exhausted, then coalitions with 2 or $M - 2$

Layer	Coalition			
1	1	0	0	0
1	0	1	1	1
2	1	1	0	0
2	0	0	1	1

Table 1: Examples of coalitions and their corresponding layers for $M = 4$. The first coalition is in Layer 1, as it contains only one feature. The second coalition is still in Layer 1 as it lacks only one feature. In the last two cases, two features are present or absent, hence the instances belong to Layer 2.

Layer	Size	Kernel SHAP	ST-SHAP (Our method)
1	30	30	30
2	210	210	210
3	910		910
4	2730	960	50
5	6006		
6	10010		0
7	12870		

Table 2: Kernel SHAP and ST-SHAP (our method) neighbor selection process for $M = 15$ and budget = 1200. The column “Size” is the number of coalitions in the layer, whereas the columns “Kernel SHAP” and “ST-SHAP” denote the number of coalitions from a layer materialized by the two different generation processes.

features included are generated, and so on. Based on the allocated budget, a layer may be either fully enumerated or not. A layer is considered **complete** if all possible coalitions within that layer have been enumerated. Conversely, if there are coalitions from a layer that have not been generated, we will refer to it as an **incomplete** layer. When focusing on a specific layer, there are two options: either fully explore that layer and then move on to the next layer, or switch to random sampling mode on all the other unexplored layers. For a layer to be fully explored, two conditions must be met:

1. The remaining budget must be sufficient to generate all the instances of that layer.
2. The weight of the layer (sum of the weights of all instances within the layer) must be significant enough to justify its exploration. Specifically, the weighted probability of drawing an instance of that layer must be greater than the probability of randomly selecting any instance.

The third column in Table 2 provides an example of SHAP’s layer-wise exploration for a budget of 1200 coalitions, and illustrates how Kernel SHAP switches to random selection after layer 2, despite having enough budget to fully explore layer 3. The fourth column of the table will be explained in the *Improving Kernel SHAP’s Stability* section.

Finally, similarly to LIME, SHAP offers the option to

control the complexity of the explanation. This can be achieved by specifying the number of features to include, i.e. the number of non-zero coefficients provided by the linear explainer. This step ensures the *interpretability* of the explanation (Ribeiro, Singh, and Guestrin 2016; Poursabzi-Sangdeh et al. 2021).

To summarize, Kernel SHAP is a method for approximating the SHAP values that relies on a weighted linear regression. The regression coefficients are estimates of the SHAP values. The weighted linear model is learned from a sample of the coalitions. The sampling procedure prioritizes lower-level layers whenever possible, and else draws randomly coalitions from higher-level layers. For each coalition, the prediction is obtained using a black-box predictive function f_x , and the weight is computed using the Shapley Kernel (Equation (6)). Finally, the weighted linear model is fitted, according to Equation (5), using the weighted coalitions and their corresponding predictions.

Reason for Instability

In this section, we illustrate the instability of Kernel SHAP with an example shown in Figure 1 and present an hypothesis on the reason of this instability. In this figure, we use an instance of the Dry Bean dataset¹ and compute three explanations on the same instance with identical parameters. We note that the four relevant features are not the same across these three executions.

To quantify stability, we compute explanations for the same instance multiple times under the same configuration and compute the Jaccard coefficient of the sets of features with non-zero coefficients in the explanations. If S_1, \dots, S_n denote those sets, this corresponds to:

$$J(S_1, S_2, \dots, S_n) = \frac{|S_1 \cap S_2 \cap \dots \cap S_n|}{|S_1 \cup S_2 \cup \dots \cup S_n|} \quad (7)$$

The coefficient ranges between 0 and 1, and values closer to 1 denote better stability.

In our previous example, the Jaccard coefficient obtained for 20 explanations (on the same instance) is 0.74 (with a budget set to 200). Further results on the instability of Kernel SHAP are presented in the *Experiments* section.

We believe that the main source of instability in Kernel SHAP explanations lies in the way coalitions are generated. As seen in the previous section, Kernel SHAP’s neighbor selection strategy introduces a lot of randomness in the process. The generation of coalitions used to fit the linear model involves layer-wise filling whenever possible, followed by random sampling (see Table 2). In practice only the lower-level layers are fully explored, whereas the higher layers are (sub-)sampled from a large and diverse population based on the remaining budget. Those samples can be very different across multiple executions of Kernel SHAP on the same instance. We believe that this stochastic process is the primary source of Kernel SHAP’s instability.

¹Dry Bean Dataset. <https://doi.org/10.24432/C50S4B>. Accessed: 2024-02-01.

Improving Kernel SHAP’s Stability

Our goal is to reduce the variability in explanations across multiple executions of Kernel SHAP on the same instance, and if possible, to achieve complete determinism. To do so, we propose to drop Kernel SHAP’s condition 2 when filling a layer (see Section *Background*, subsection *Kernel SHAP*, condition 2 to fully explore a layer), that is, fill the layers in increasing order as long as there is still enough budget. Unlike Kernel SHAP, if the budget is not sufficient to fill a layer, random generation only takes place in that layer – instead of all the remaining layers. Then, to eliminate variability completely, we can set the budget in a way that only complete layers are considered.

Let us consider the case where $M = 15$ depicted in Table 2, which illustrates Kernel SHAP’s neighbor generation process for 7 layers. To generate elements of the first layer, we consider the cases where only one feature is included and the cases where only 1 feature is excluded. We can easily enumerate all possible coalitions within this layer, resulting in exactly $2 \cdot C_{15}^1 = 30$ coalitions. If the allocated budget is 30, we can be sure that the same 30 coalitions will be generated every time we run SHAP. In that line of thought, budgets of 30 or 240 ($30 + 210$) guarantee full determinism in SHAP because they entail complete layers.

Now, if the budget is 1200 as in the example, Kernel SHAP will saturate layers 1 and 2 (with 30 and 210 neighbors respectively), which leaves us with a remaining budget of $1200 - 210 - 30 = 960$ instances. Kernel SHAP will sample those instances from all the remaining layers, whereas we propose to saturate layer 3 with 910 neighbors and sample the remaining 50 neighbors randomly from layer 4. As our experiments show, such an approach, that we call **ST-SHAP** (*Stable SHAP*), mitigates the deleterious effects of sub-sampling in the stability of the explanations. This is so because we restrict the random sampling of neighbors to a smaller and less diverse universe, i.e., layer 4 in ST-SHAP vs. layers 3 to 7 in Kernel SHAP.

We will also show the benefits of setting the number of neighbors in a way that allows only complete layers, thereby completely eliminating variability.

Experiments

We conduct a series of experiments to verify the effectiveness of ST-SHAP at mitigating the instability of Kernel SHAP.

Experimental Protocol


Environment. The experiments are conducted on a computer with a 12th generation Intel-i7 CPU (14 cores with HT, 2.3-4.7 GHz Turbo Boost). We use the official Kernel SHAP implementation provided by the authors (Lundberg and Lee 2017). ST-SHAP² is based on this implementation. We use the scikit-learn python library³ to fit the black-box models that we aim to explain.

²<https://github.com/gwladyskelodjou/st-shap>


³<https://scikit-learn.org/>

Feature	Value
Area	55487.0
Perimeter	915.3
MajorAxisLength	359.2
MinorAxisLength	197.7
AspectRatio	1.82
Eccentricity	0.83
ConvexArea	56143.0
EquivDiameter	265.8
Extent	0.65
Solidity	0.99
Roundness	0.83
Compactness	0.74
ShapeFactor1	0.006
ShapeFactor2	0.001
ShapeFactor3	0.55
ShapeFactor4	0.99


(a) Input instance



(b) Execution 1



(c) Execution 2



(d) Execution 3

Figure 1: Illustration of the instability of SHAP on an instance from the Dry Bean dataset. The explanation size is 4, so each explanation computation yields four features with non-zero values.

Datasets. We use several datasets from the UCI Machine Learning Repository and other popular real-world datasets, namely: Boston (Harrison Jr and Rubinfeld 1978) and Movie, which are regression datasets, and Default of Credit Card Clients (Yeh 2016), Adult (Becker and Kohavi 1996)⁴, Dry Bean⁵, Spambase (Hopkins et al. 1999), HELOC⁶, and Wisconsin Diagnostic Breast Cancer (Wolberg et al. 1995), which are classification datasets. Table 3 summarizes the information about the experimental datasets.

Black-Box Models. We also use various black-box models, including Support Vector Machine (SVM), Random Forest (RF), Logistic Regression (LR), and Multi-layer Perceptron (MLP). For each dataset, we randomly select 10 instances from the test set. Finally, we compute the explanation 20 times for each explained instance to derive the corresponding Jaccard value that estimates stability.

Metrics. We evaluate stability by measuring the evolution of the Jaccard coefficient (Equation (7)) as we change the budget. We test budgets that lead to complete layers for layers 1 to 3, as well as budgets that do not: {10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 200, 500, 1000, 2000}.

⁴We use a reformatted version released by the authors of SHAP.

⁵<https://doi.org/10.24432/C50S4B>. Accessed: 2024-02-01.

⁶<https://community.fico.com/s/explainable-machine-learning-challenge?tabset-158d9=3>. Accessed: 2024-02-01.

Dataset	Size	#features	Task	Explanation size
Movie	505	19	Regression	4
Boston	506	13	Regression	4
Dry Bean	13611	16	Classification	4
Credit Card	30000	23	Classification	4
Adult	32561	12	Classification	4
Spambase	4601	57	Classification	4
HELOC	10459	23	Classification	4
WDBC	569	30	Classification	4

Table 3: Summary of datasets used in our experiments. The *Explanation size* indicates the number of non-zero coefficients returned by the linear model of the explanation. These are the most significant features associated to a particular prediction (the absolute value of each coefficient indicates the magnitude of that importance).

Results

For the sake of conciseness, we refer to Kernel SHAP as SHAP in all the charts. Figure 2a presents the results obtained on the Boston dataset. We report the average Jaccard coefficient among the 10 instances for each of the studied budgets. Layers 1, 2 and 3 become complete when the budgets are 26, 182 and 754 respectively. We can see that, for




Figure 2: Evolution of the Jaccard coefficient for SHAP and ST-SHAP on Boston, Dry Bean, HELOC, and Spambase datasets. The vertical lines represent the budgets that result in a complete layer.

ST-SHAP, the Jaccard coefficient equals 1.0 when the budget corresponds to complete layers. This indicates a complete absence of variability in the process to compute the explanation. We also notice that ST-SHAP’s stability outperforms Kernel SHAP’s as we approach budgets with complete layers. In addition to these observations, we highlight in Figure 2b with the Dry Bean dataset that ST-SHAP is generally more stable than Kernel SHAP, even on incomplete layers. Figure 2d showcases results for Spambase dataset, a larger dataset. For complete layers, only layers 1 and 2 have been calculated, as layer 3 requires a budget of 61826 coalitions. Budgets for incomplete layers remain unchanged, except with the addition of a budget of 3500 coalitions. We observe the same trends as in the previous datasets. The results presented here cover one regression and three classification datasets. The results for the other datasets can be found in (Kelodjou et al. 2023, Appendix 2). Similar observations apply to these datasets.

Since stable explanations of low fidelity are as useless as unstable explanations, we examine the impact of our method on the fidelity of the explanations. To assess fidelity, we compare the outputs of the black-box model with those of the Kernel SHAP and ST-SHAP linear approximation models. For this purpose, we use the R^2 statistic for regression tasks (using the Boston and the Movie datasets) and the *accuracy* for classification tasks (for other datasets). In both cases the metrics assess the quality of the linear approxima-

tion provided by Kernel SHAP or ST-SHAP at “adhering” to the black box’s outcomes – so the higher they are the better. This form of fidelity is referred to as *adherence* in the literature (Visani, Bagli, and Chesani 2020; Gaudel et al. 2022).

Figure 3a is produced by taking the average R^2 -score obtained for each budget. Each R^2 -score is obtained by comparing the results of Kernel SHAP (resp. ST-SHAP) on the set of coalitions used to train the linear explainer. From this figure, it can be observed that the fidelity of ST-SHAP is comparable to Kernel SHAP’s. It is also noteworthy that the fidelity is generally high, even for small budgets (e.g., by considering only the neighbors in Layer 1). Figure 3b illustrates the fidelity comparison between Kernel SHAP and ST-SHAP on the Dry Bean dataset. The results demonstrate a generally good fidelity for both methods, with ST-SHAP sometimes outperforming Kernel SHAP. Furthermore, for both methods, fidelity tends to decrease with an increasing budget. This trend can be attributed to the fact that, as the budget increases, we sample neighbors from upper layers, which confronts the linear regression with a more diverse set of coalitions. Depending on the black-box, this diversity can make the optimization problem more challenging, resulting in lower fidelity compared to a regression trained on the first layers. In Figure 3d, where we present the results for the Spambase dataset, it can be observed that ST-SHAP generally exhibits better fidelity than Kernel SHAP.

The results for the remaining datasets are provided in

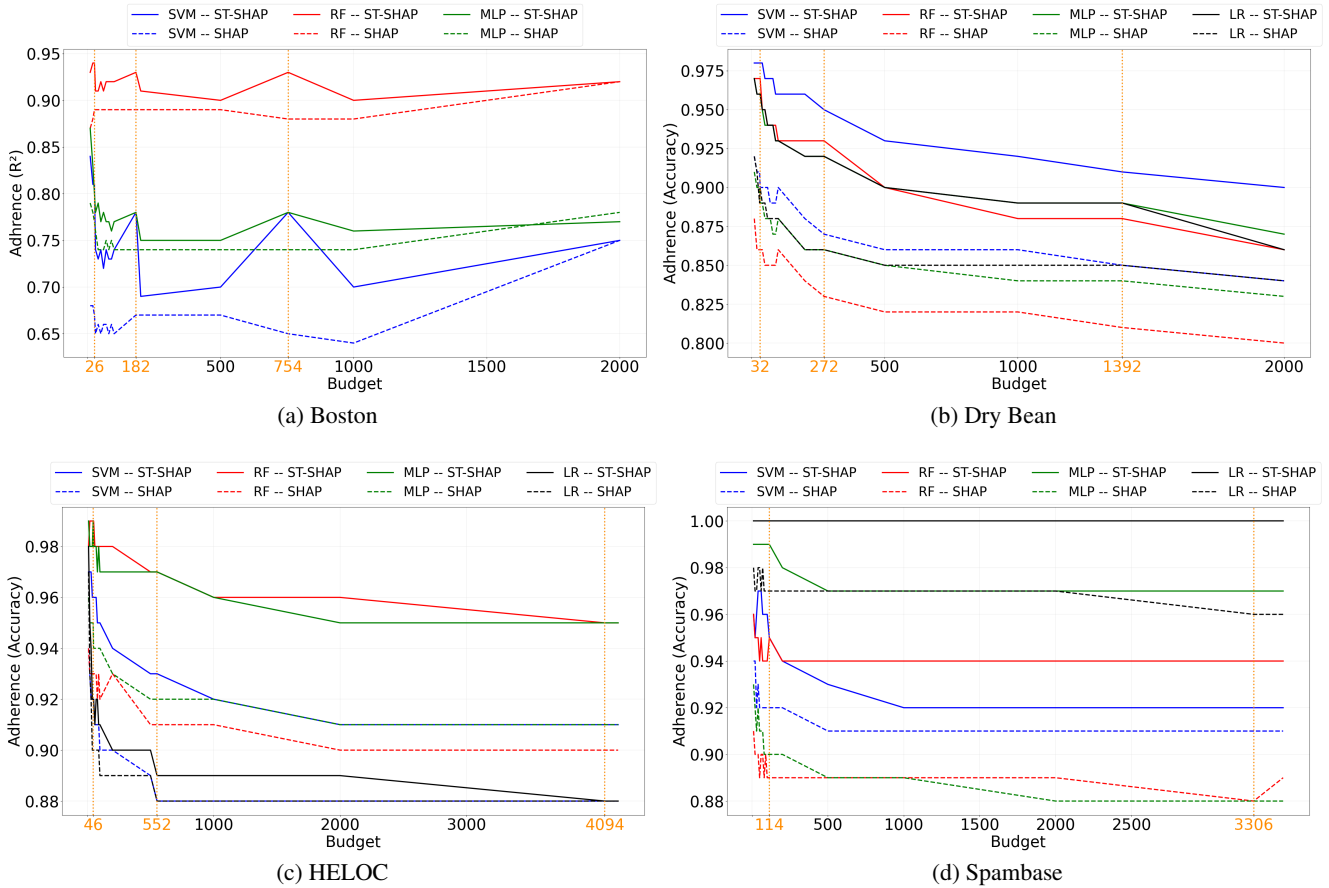


Figure 3: Adherence (R^2 -score for Boston and Accuracy for others) vs. budget size for Boston, Dry Bean, HELOC, and Spambase datasets.

(Kelodjou et al. 2023, Appendix 2). These datasets yield the same conclusions. Additional experiments comparing the exact SHAP values with the coefficients obtained with ST-SHAP show that the approximation of the SHAP values provided by ST-SHAP is as good as that of Kernel SHAP (Kelodjou et al. 2023, Appendix 2).

Our experiments on stability and fidelity lead us to conclude that our layer-wise strategy to complete layers of coalitions in Kernel SHAP reduces variability satisfactorily, leading to stability. This stability is maximal (Jaccard score equals 1.0) when the budget guarantees fully complete layers. Moreover, our methods do not have an impact on the adherence of the explanations. This means that users can now choose which layers to use depending on the available time and computing resources, since a larger number of coalitions translates into longer execution times. After examining the stability and fidelity results, we note that learning a surrogate on the coalitions of **Layer 1** not only offers complete stability and interesting fidelity, but it is also very fast to compute because there are only $2M$ of such coalitions (where M is the total number of features). This makes Layer 1 a potentially interesting choice when learning attribution scores. In the following, we conduct a theoretical study of the attribution scores obtained with such a strategy.

First Layer Attribution Analysis

In this section, we investigate the attribution values obtained using exclusively the coalitions from Layer 1. These attribution values are of particular interest because they are the cheapest to compute and guarantee stable explanations. We first characterize these attribution values, and demonstrate their good theoretical properties, then we show to which extent they align with the exact SHAP values on real datasets, and finally we compare the execution times.

Theoretical Analysis

In order to simplify the notations, we denote by $N = \{1, \dots, M\}$ the set of all features and we define f on $\mathcal{P}(N)$ as shown by Equation (8). $\mathbb{1}_S \in \{0, 1\}^M$ is defined by $(\mathbb{1}_S)_i = 1$ if $i \in S$ and 0 otherwise, with S being a coalition of features:

$$f : \mathcal{P}(N) \rightarrow \mathbb{R} \\ S \mapsto f_x(\mathbb{1}_S). \quad (8)$$

Coalitions of Layer 1 are of two types: coalitions with a single feature present (a single 1) and coalitions with a single feature absent (a single 0). In terms of marginal contribution, this amounts to averaging the individual contribution: $f(\{j\}) - f(\emptyset)$, and the effect of removing the feature from the input: $f(N) - f(N \setminus \{j\})$, $\forall j \in N$.

By also satisfying the *local accuracy* constraint (Lundberg and Lee 2017), i.e., $\sum_{i=1}^M \phi_i = f(x) - f(\emptyset)$, we obtain Theorem 1, which expresses a closed form for the attribution values of ST-SHAP when restricted to Layer 1.

Theorem 1 (Attribution values with Layer 1). *For any feature j , the attribution value ϕ_j computed by ST-SHAP when filling exactly Layer 1 is:*

$$\phi_j = \tilde{\phi}_j + \frac{1}{M} \left(f(N) - f(\emptyset) - \sum_{i=1}^M \tilde{\phi}_i \right) \quad (9)$$

where for any i , $\tilde{\phi}_i = \frac{f(\{i\}) - f(\emptyset) + f(N) - f(N \setminus \{i\})}{2}$.

Proof. We provide below a summary of the main steps of the proof, with some intermediate steps omitted. For a detailed proof of this theorem refer to Kelodjou et al. (2023, Appendix 1).

Determining the attribution values ϕ_j when restricting to coalitions of layer 1 is equivalent to solving the optimization problem formulated in Equation (10). The weighting kernel π_x w.r.t Equation (5) is omitted because all coalitions in layer 1 have the same weight. The formulation of the minimization function then becomes:

$$\arg \min_{\phi} \sum_{z' \in \{0,1\}^M} (g(z') - f_x(z'))^2. \quad (10)$$

Coalitions z' of layer 1 are characterized by $|z'| = 1$ or $|z'| = M - 1$.

- When $|z'| = 1$ with $j \in N$ being the only present feature, we obtain the equation:

$$g(z') = \phi_0 + \phi_j.$$

We aim to determine g such that $g(z')$ provides a close approximation to $f(\{j\})$, i.e., $g(z') \approx f_x(z')$ where $f_x(z') = f(\{j\})$.

- When $|z'| = M - 1$, with $j \in N$ being the only absent feature, we get:

$$g(z') = \phi_0 + \sum_{\substack{i=1 \\ i \neq j}}^M \phi_i.$$

We aim to determine g such that $g(z')$ provides a close approximation to $f(N \setminus \{j\})$, i.e., $g(z') \approx f_x(z')$ where $f_x(z') = f(N \setminus \{j\})$.

Given the following terms:

- $f(\{j\}) - f(\emptyset)$ denoted by $\Delta_{j\emptyset}$
- $f(N \setminus \{j\}) - f(\emptyset)$ denoted by $\Delta_{j\bar{0}}$
- $f(N) - f(N \setminus \{j\})$ denoted by $\Delta_{N\bar{j}}$
- $f(N) - f(\emptyset)$ denoted by Δ ,

this is tantamount to solve a minimization problem on the ϕ_j scores with two constraints: local accuracy ($\sum_{i=0}^M \phi_i = f(N)$), and missingness ($\phi_0 = f(\emptyset)$). The minimization problem is defined by:

$$\arg \min_{\phi} \mathcal{L}(\phi) \quad (11)$$

where $\mathcal{L}(\phi)$ is defined by:

$$\mathcal{L}(\phi) = \sum_{j=1}^M (\phi_j - \Delta_{j\emptyset})^2 + \sum_{j=1}^M \left(\left(\sum_{\substack{i=1 \\ i \neq j}}^M \phi_i \right) - \Delta_{j\bar{0}} \right)^2 \quad (12)$$

$$= \sum_{j=1}^M (2\phi_j^2 - 2(\Delta_{N\bar{j}} + \Delta_{j\emptyset})\phi_j + \Delta_{N\bar{j}}^2 + \Delta_{j\emptyset}^2). \quad (13)$$

Because of the additional constraints (local accuracy), solving this problem requires using the method of the Lagrange multiplier, which leads us the following alternative optimization problem:

$$\arg \min_{\phi, \lambda} \mathcal{L}'(\phi, \lambda) = \mathcal{L}(\phi) + \lambda \left(\sum_{j=1}^M \phi_j - \Delta \right). \quad (14)$$

This function is minimized when

$$\frac{\partial \mathcal{L}'}{\partial \phi} = 0 \quad \text{and} \quad \frac{\partial \mathcal{L}'}{\partial \lambda} = 0.$$

For a given ϕ_j the new equations are:

$$\frac{\partial \mathcal{L}'}{\partial \phi_j} = 0 = 4\phi_j - 2\Delta_{N\bar{j}} - 2\Delta_{j\emptyset} + \lambda \quad (15)$$

$$\frac{\partial \mathcal{L}'}{\partial \lambda} = 0 = \Delta - \sum_{j=1}^M \phi_j. \quad (16)$$

Adding up the M variants of Equation (15), we can solve for λ :

$$\lambda = \frac{\sum_{j=1}^M (2\Delta_{N\bar{j}} + 2\Delta_{j\emptyset}) - 4 \sum_{j=1}^M \phi_j}{M}. \quad (17)$$

We can now plug Equation (17) into Equation (15), leading to our definition of ϕ_j for any $j \in N$:

$$\phi_j = \tilde{\phi}_j + \frac{1}{M} \left(f(N) - f(\emptyset) - \sum_{i=1}^M \tilde{\phi}_i \right). \quad (18)$$

□

It can be noted that with such a formula, the three properties: Linearity, Symmetry, and Efficiency are verified (Kelodjou et al. 2023, Appendix 1). Therefore, the proposed attribution method belongs to the LES family of attribution scores (Ruiz, Valenciano, and Zarzuelo 1998; Chameni Nembua and Andjiga 2008; Radzik and Driessen 2013), which also includes the Shapley values. LES values are based on marginal contributions, providing feature contributions and interpretations very close to the Shapley values (Condevaux, Harispe, and Mussard 2022).

Layer 1 versus Shapley Values

We observed that the attribution values derived solely from the coalitions within Layer 1 are more straightforward to calculate and fulfill the three axioms of the LES family of attribution scores, also called LES values (Ruiz, Valenciano, and Zarzuelo 1998; Condevaux, Harispe, and Mussard 2022). We now compare these values with the actual SHAP values (computed by materializing all coalitions) using various datasets. This comparison is based on the following considerations:

- We include all the features in the explanation, meaning that the explanation size is M , i.e., the number of features in the target dataset.
- We use the Kendall rank correlation coefficient to compare the rankings of the features within the explanation. This coefficient ranges from -1 to 1. Values close to 1 indicate that the Layer 1’s attribution ranks are in agreement with the actual SHAP values.
- To compare the disparity between the actual SHAP values and the Layer 1 attribution scores, we use the coefficient of determination, R^2 . While mainly used to evaluate wellness of fit for machine learning models, we remark that it is suitable in our context. If we consider the SHAP values as the “real” values and the Layer 1 scores as the “predicted” values designed to fit the real values, the R^2 measures the proportion of the variance of the real SHAP values “explained” by the Layer 1 attribution scores. When the R^2 coefficient is 1 then both sets of scores agree, whereas values closer to 0 suggest that our approximation is as good as the naive attribution that assigns the same value to all features. Values smaller than 0 denote a performance below the naive baseline.

The experiments of this section are conducted on a subset of the datasets discussed in the *Experiments* section, where computing the exact SHAP values is feasible. For black-box models, we employ the same models as described in the *Experiments* section. The experiments are conducted on the whole test set (except for Credit card dataset, in which only 10 random instances on the test set are used, due to high execution time required to compute the actual SHAP values) and we report averages and medians. Table 4 reports the corresponding results. We observe from Kendall’s coefficient that the features ranks are almost the same between Layer 1’s attribution scores and the exact SHAP values. The magnitudes of those scores also remain very similar as shown by the R^2 scores in a majority of cases. This makes Layer 1 attribution scores a very appealing choice for very large datasets: they are stable and very fast to compute, while still agreeing with the exact SHAP values.

Execution Times

The runtime of ST-SHAP Layer 1 is significantly smaller, compared to the computation of exact SHAP values, with complexities of $O(M)$ and $O(2^M)$ respectively.

Considering that the standard Kernel SHAP is recommended with a budget of 2000 ($2 \cdot M + 2^{11}$), utilizing ST-SHAP Layer 1 shows to be up to one order of magnitude

		Kendall τ			R^2 -Score		
		SVM	RF	MLP	SVM	RF	MLP
Boston	μ	0.95	0.91	0.959	0.98	0.99	0.99
	M	0.97	0.92	0.97	0.99	0.99	0.99
Adult	μ	0.7	0.68	0.87	1.0	0.65	0.41
	M	0.6	0.7	0.9	1.0	0.77	0.69
Dry Bean	μ	0.86	0.75	0.84	-0.07	0.74	0.75
	M	0.9	0.76	0.88	0.51	0.79	0.8
Movie	μ	1.0	0.89	0.85	1.0	0.99	0.95
	M	1.0	0.96	0.86	1.0	0.99	0.95
Credit Card	μ	0.84	0.64	0.79	0.94	0.02	0.81
	M	0.89	0.71	0.81	0.98	0.21	0.83

Table 4: Kendall’s τ and R^2 -Score between the SHAP values and the attribution scores of ST-SHAP from the coalitions of Layer 1. μ : Mean, M: Median.

faster than Kernel SHAP⁷. This is so because the budget is significantly lower than 2000 as outlined in (Kelodjou et al. 2023, Appendix 2).

Conclusion

In this paper, we have investigated the instability issues of the Kernel SHAP estimator and proposed a novel neighbor selection approach that achieves full stability. Our experimental results demonstrate that our approach does not compromise fidelity. We have also conducted a theoretical analysis of the coefficients obtained by applying ST-SHAP on the neighbors of Layer 1. We define those scores formally and show experimentally that they remain very close to the actual SHAP values. This makes this approach an attribution method in itself, that incurs faster computation while retaining most of the desirable properties of attribution scores.

In the future we are interested in understanding the properties of black-box models that guarantee good approximations of the SHAP values when trained on subsets of the coalitions space. This could include the definition of other attribution methods, and the relationship between the budget and the complexity of the target black-box we aim to explain.

References

- Alvarez-Melis, D.; and Jaakkola, T. S. 2018. On the Robustness of Interpretability Methods. In *Proceedings of the 2018 ICML Workshop in Human Interpretability in Machine Learning*, 66–71. Stockholm, Sweden: ICML.
- Becker, B.; and Kohavi, R. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>. Accessed: 2024-02-01.
- Chameni Nembua, C.; and Andjiga, N. G. 2008. Linear, efficient and symmetric values for TU-games. *Economics Bulletin*, 3(71): 1–10.

⁷It is also up to three orders of magnitude faster than computing the exact SHAP values.

- Condevaux, C.; Harispe, S.; and Mussard, S. 2022. Fair and Efficient Alternatives to Shapley-based Attribution Methods. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 309–324. Springer.
- Gaudel, R.; Galárraga, L.; Delaunay, J.; Rozé, L.; and Bhargava, V. 2022. s-LIME: Reconciling locality and fidelity in linear explanations. In *International Symposium on Intelligent Data Analysis*, 102–114. Springer.
- Ghorbani, A.; and Zou, J. 2019. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*, 2242–2251. PMLR.
- Harrison Jr, D.; and Rubinfeld, D. L. 1978. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 5(1): 81–102.
- Hopkins, M.; Reeber, E.; Forman, G.; and Suermondt, J. 1999. Spambase. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C53G6X>. Accessed: 2024-02-01.
- Kelodjou, G.; Roze, L.; Masson, V.; Galarraga, L.; Gaudel, R.; Tchente, M.; and Termier, A. 2023. Shaping Up SHAP: Enhancing Stability through Layer-Wise Neighbor Selection. arXiv:2312.12115.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Poursabzi-Sangdeh, F.; Goldstein, D. G.; Hofman, J. M.; Wortman Vaughan, J. W.; and Wallach, H. 2021. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.
- Radzik, T.; and Driessen, T. 2013. On a family of values for TU-games generalizing the Shapley value. *Mathematical Social Sciences*, 65(2): 105–111.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- Ruiz, L. M.; Valenciano, F.; and Zarzuelo, J. M. 1998. The family of least square values for transferable utility games. *Games and Economic Behavior*, 24(1-2): 109–130.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Shapley, L. S.; et al. 1953. A Value for n-Person Games. In Kuhn, H. W.; and Tucker, A. W., eds., *Contributions to the Theory of Games (AM-28), Volume II*, 307–318. Princeton University Press, Princeton. ISBN 978-1-4008-8197-0.
- Song, E.; Nelson, B. L.; and Staum, J. 2016. Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1): 1060–1083.
- Strumbelj, E.; and Kononenko, I. 2010. An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, 11: 1–18.
- Visani, G.; Bagli, E.; and Chesani, F. 2020. OptiLIME: Optimized LIME Explanations for Diagnostic Computer Algorithms. In *Proceedings of the CIKM 2020 Workshops co-located with 29th ACM International Conference on Information and Knowledge Management*, volume 2699 of *CEUR Workshop Proceedings*.
- Visani, G.; Bagli, E.; Chesani, F.; Poluzzi, A.; and Capuzzo, D. 2022. Statistical stability indices for LIME: Obtaining reliable explanations for machine learning models. *Journal of the Operational Research Society*, 73(1): 91–101.
- Wolberg, W.; Mangasarian, O.; Street, N.; and Street, W. 1995. Breast Cancer Wisconsin (Diagnostic). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5DW2B>. Accessed: 2024-02-01.
- Yeh, I.-C. 2016. default of credit card clients. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C55S3H>. Accessed: 2024-02-01.
- Zhou, Z.; Hooker, G.; and Wang, F. 2021. S-lime: Stabilized-lime for model explanation. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2429–2438.