

A Multilinear Sampling Algorithm to Estimate Shapley Values

Ramin Okhrati
University College London
London, United Kingdom
r.okhrati@ucl.ac.uk

Aldo Lipani
University College London
London, United Kingdom
aldo.lipani@ucl.ac.uk

Abstract—Shapley values are great analytical tools in game theory to measure the importance of a player in a game. Due to their axiomatic and desirable properties such as efficiency, they have become popular for feature importance analysis in data science and machine learning. However, the time complexity to compute Shapley values based on the original formula is exponential, and as the number of features increases, this becomes infeasible. Castro et al. [1] developed a sampling algorithm, to estimate Shapley values. In this work, we propose a new sampling method based on a multilinear extension technique as applied in game theory. The aim is to provide a more efficient (sampling) method for estimating Shapley values. Our method is applicable to any machine learning model, in particular for either multi-class classifications or regression problems. We apply the method to estimate Shapley values for multilayer perceptrons (MLPs) and through experimentation on two datasets, we demonstrate that our method provides more accurate estimations of the Shapley values by reducing the variance of the sampling statistics.

Index Terms—shapley values, feature attribution methods, sampling

I. INTRODUCTION

Explaining models have becoming an important aspect of data science, in particular, with the growing complexity of machine learning models. For motivations, discussions, and different explainability models under a unified framework, we refer to the work of Lundberg and Lee [2] where the methods LIME [3], DeepLIFT [4], Layer-Wise Relevance Propagation [5], and exact Shapley values [6] together with their SHAP methods are explained. The main reason for the popularity of Shapley values is that they satisfy certain interesting axioms (which we will discuss in Section III), most notably the efficiency axiom, that are desirable in analyzing interpretability and attributions.

It is worth noting that Shapley values are a local interpretability attribution method in the sense that they quantify the explanation of a specific model by assigning a value to each input feature. This value can be then interpreted as the contribution (or relevance) of that feature towards the output of the model. Although several non-axiomatic methods are developed in particular for neural networks, an advantage of an axiomatic approach, such as Shapley values, is to have a profound theoretical background that mitigates the risk of misleading or unreliable interpretations. While there

are other attribution methods which have their own merits, we specifically focus on estimating Shapley values through sampling methods.

The main drawback of the calculation of Shapley values using the original definition (aka *exact Shapely values*) is its exponential time complexity. We divide Shapley values estimation methods into two main categories, methods that do not converge towards the exact Shapely values, and methods that do converge. In the first category, we have two types of methods. The first type are semi-closed form solution methods to approximate Shapley values which provide a decent approximation for simple games (see the work of Castro et al. [1] where they make a comparison of this method with their sampling approach). These solutions are normally based on the central limit theorem (see the work of Owen [7]) which are computationally cheap. The second type of the first category are data-driven methods such as regression and linear based techniques [8, 9], quantitative input influence approaches [10], and DASP [11]. The second category includes the Shapley sampling method developed by Castro et al. [1] which in general becomes computationally expensive as the number of feature increases, however, this has the advantage of converging to the exact Shapley values.

Our method falls within the second category in which we provide a sampling type algorithm to estimate Shapley values based on a multilinear extension technique as applied in game theory [7], which converges to the exact Shapley values. As mentioned earlier, semi-closed form solutions have already been obtained from this multilinear extension techniques for simple games [7], however, to the best of our knowledge, this is the first time that a sampling type algorithm is extracted from the multilinear extension techniques. Furthermore, our algorithm is applicable for calculating the Shapley values of features of any machine learning model. In particular, we apply the algorithm for the case of MLPs in classification problems.

The sampling methods such as the one developed by Castro et al. [1] are used as a ground truth for Shapley value estimations or attributions, see for instance the work of Ancona et al. [11]. Their experiments show that for up to a certain number of network evaluations, their method provides a more accurate solution. In this paper we compare against the sampling algorithm developed by Castro et al. [1] and from now on we will refer to this algorithm as *Castro sampling*.

As our experiments show, comparing to the Castro sampling, our method reduces the variance of the estimator which results in a more accurate and time-efficient algorithm. The goal is to present our algorithm in its original form though it could be improved in several directions which we leave to future work. In our experimental analysis, first, we use a dataset with a small number of features in which we can calculate the exact Shapley values. Knowing the exact Shapley values in this case, we can then compare the quality of our method, in terms of both accuracy and time-efficiency against Castro sampling. As the results confirm, the error of our approximation is much lower. Next, we then apply the method also on a dataset with a larger number of features. While it is not possible to obtain the exact Shapley values in this case, we can perform a variance analysis to show that our estimator admits a lower variance compared to that of Castro sampling.

The remainder of this paper goes as follows. In § II, we discuss related works. In § III, we introduce the background and preliminaries necessary to explain our method in § IV. § V is devoted to the experimental setup. Results are presented and discussed in § VI. We conclude in § VII.

II. RELATED WORK

Multilinear extension methods have been already applied in calculating feature attributions. Jones and Wilson [12] present a multilinear extension technique, which is then applied to multichoice games in the context of Shapley values using the central limit theorem. Michalak et al. [13] present analytical formulas for Shapley value-based centrality in both weighted and unweighted networks algorithms. The base of our method is the multilinear extension technique as developed by Owen [7]. This work is further discussed in the next section. Despite other multilinear extension techniques applied to estimate Shapley values, our method does indeed converge to the exact Shapley values. Although sampling methods should converge to the exact values, the sampling nature still remains there, and it becomes computationally cumbersome in the presence of a large number of features. So the efficiency and applicability of the method boils down to its convergence speed.

Depending on the structure of the machine learning model, in certain cases, simpler approximations might be possible. Fatima et al. [9] introduce a polynomial time approximation method to estimate Shapley values for weighted voting games. A similar idea is also used by Ancona et al. [11] for explaining deep neural networks where in contrast to the sampling method developed by Castro et al. [1], their algorithm does not converge to the exact Shapley values, however, for up to a fixed number of instructions, their algorithm provides a more accurate estimation.

It is worth noting that there are other attribution methods as well. Lundberg and Lee [2] introduce the KernelSHAP method, which combines sampling with lasso regression. There are also Integrated Gradient [14], and DeepLIFT and its two flavors Rescale and RevealCancel [15, 4]. However, among all these methods, only KernelSHAP (without regularization) converges to the exact Shapley values.

III. BACKGROUND AND PRELIMINARIES

A. Feature Attribution Through Shapely Values

Consider a machine learning model ν from \mathbb{R}^{n+1} into \mathbb{R} with inputs represented by $X = (x_0, x_1, \dots, x_n)$ in the feature space \mathbb{R}^{n+1} , where $x_0 = 1$ is to account for the bias term, and $n \geq 1$ is a positive integer.

A prime question in many applied fields such as data science and finance is to provide an algorithm that measures the contribution (also called attribution or importance) of a given feature x_j , $j = 0, 1, \dots, n$, in the output $\nu(X)$ where X is the input of the machine learning model. But how should such notion of importance be defined?

Let us have a short excursion to game theory where the notion of contribution is well defined for games. In game theory language, a game is a real-valued function ξ defined on the player set (called features in our setup) $\{x_0, x_1, \dots, x_n\}$, i.e., $\xi : 2^{\{x_0, x_1, \dots, x_n\}} \rightarrow \mathbb{R}$. In contrast to machine learning where the inputs are tuples (vectors) of features, in game theory the inputs are set of players.

Suppose that we want to define a notion of contribution for the player x_j , $j = 0, 1, \dots, n$, of a game expressed by ξ . Let us denote the contribution of feature x_j by $\Psi_j(\xi)$. It is argued by Shapley [6] that $\Psi_j(\xi)$ is a reasonable notion of contribution, if it satisfies (at least) the following axioms:

Efficiency: $\sum_{j=0}^n \Psi_j(\xi) = \xi(X)$. That is, the sum of all contributions is equal to the output of the game.

Symmetry: For any two players x_i and x_j , if $\xi(A \cup \{x_i\}) = \xi(A \cup \{x_j\})$ for all subset A (of the player set) that do not contain neither x_i nor x_j , then $\Psi_j(\xi) = \Psi_i(\xi)$. That is, if two players have the same influence on any coalition their contributions should be the same.

Linearity: For two games ξ and ω with the same player set and for all players x_j and a real number α , we have $\Psi_j(\xi + \omega) = \Psi_j(\xi) + \Psi_j(\omega)$ and $\Psi_j(\alpha\xi) = \alpha\Psi_j(\xi)$. That is, when two independent games are combined, the total contribution of a player is equal to the sum of individual contributions on each game.

Null player (feature): For a player x_j and a game ξ , if $\xi(A \cup \{x_j\}) = \xi(A)$ for all subsets A (of the player set) that do not contain x_j (such player is called null player) then $\Psi_j(\xi) = 0$.

Remarkably, it is proved by Shapley [6] that $\Psi_j(\xi)$ must satisfy the following representation:

$$\Psi_j(\xi) = \sum_{A \subseteq X \setminus \{x_j\}} \frac{|A|! (n - |A|)!}{(n + 1)!} (\xi(A \cup \{x_j\}) - \xi(A)), \quad (1)$$

where $|A|$ is the cardinality of the set A , i.e., the number of elements in A . In other words, the only measurement that satisfies the four axioms is given by Equation (1), which is called the Shapley value (see the work of Moulin [16] for more explanations).

Given the desirable axioms, it appears that Shapley values are suitable measurements of a player's contribution in a

game which make them also suitable candidates for measuring contributions of features in machine learning. Concretely, a machine learning model ν on a feature space (player set in game theory) can be considered as a game, and a tuple of features $X = (x_0, x_1, \dots, x_n)$ is interpreted as the set of players $\{x_0, x_1, \dots, x_n\}$. Vice versa, for a set of players (each player is interpreted as a feature), we can create a tuple where a missing feature is substituted by zero. The last point is related to the concept of baseline in machine learning where we substitute a missing feature with a zero baseline. Other choices of baselines are also available [17], but this does not affect our methodology.

From the above argument, we can formally define the contribution of a feature x_j in a machine learning model ν through Shapley value as follows:

Definition III.1. *Given a feature x_j and a machine learning model ν , we define its contribution towards $\nu(X)$ as the Shapley value of the feature x_j for ν , that is:*

$$S_j(\nu) = \sum_{A \subseteq X \setminus \{x_j\}} \frac{|A|! (n - |A|)!}{(n + 1)!} (\nu(A \cup \{x_j\}) - \nu(A)),$$

where $|A|$ is the cardinality of the set A , and with some abuse of notation, $\nu(A \cup \{x_j\})$ and $\nu(A)$ must be understood as the evaluation of ν for the corresponding tuples obtained respectively from $A \cup \{x_j\}$ and A , through replacing a missing feature by zero in the tuples.

B. Multilinear Extensions Method

Next, we provide a brief review of the multilinear extension method as developed by Owen [7]. In order to do so, first, we need to impose a probabilistic structure on the feature space in the following sense. For a fixed $j = 0, 1, \dots, n$ and a random subset E_j of $\{x_0, x_1, \dots, x_n\} \setminus \{x_j\}$, the probability of any feature (except x_j) being in E_j is equal to q where $0 \leq q \leq 1$; more precisely, this means that:

$$\mathbb{P}(\{x_k \in E_0\}) = q, \text{ for all } k = 1, 2, \dots, n,$$

$$\mathbb{P}(\{x_k \in E_n\}) = q, \text{ for all } k = 0, 1, \dots, n - 1,$$

and for $j = 1, 2, \dots, n - 1$,

$$\mathbb{P}(\{x_k \in E_j\}) = q, \text{ for all } k = 0, 1, \dots, j - 1, j + 1, \dots, n.$$

Furthermore, we assume that sampling is done so that the events $\{x_k \in E_j\}_{k \neq j, 0 \leq k \leq n}$ are mutually independent. Note that because E_j is a random subset, $\nu(E_j)$ becomes a random variable.

Given the above probabilistic structure, a probabilistic representation of the Shapley values is provided as follows [7]:

$$S_j(\nu) = \int_0^1 e_j(q) dq, \quad (2)$$

where

$$e_j(q) = \mathbb{E}[\nu(E_j \cup \{x_j\}) - \nu(E_j)]. \quad (3)$$

The integral of Equation (2) can be interpreted as either Riemann or Lebesgue integral. Equations (2) and (3) are the

Algorithm 1 Owen Sampling

Input: $\nu, X = (x_0, x_1, \dots, x_n), Q, M$

Output: $\hat{S}(\nu) = (\hat{S}_0(\nu), \hat{S}_1(\nu), \dots, \hat{S}_n(\nu))$

```

1:  $\hat{S}_j(\nu) := 0$ , for all  $j$ 
2: for  $q = 0, 1/Q, 2/Q, \dots, 1$  do
3:    $e_j := 0$ , for all  $j$ 
4:   for  $m = 1, 2, \dots, M$  do
5:      $I_m^{(q)} \leftarrow (b_j : b_j \sim \text{Bern}(q), \text{ for all } j)$ 
6:      $h_{m,j}^{(q)} \leftarrow \nu(I_m^{(q)} \odot X + X^{(j)}) - \nu(I_m^{(q)} \odot X)$ , for all  $j$ 
7:      $e_j \leftarrow e_j + h_{m,j}^{(q)}$ , for all  $j$ 
8:   end for
9:    $\hat{S}_j(\nu) \leftarrow \hat{S}_j(\nu) + e_j$ , for all  $j$ 
10: end for
11:  $\hat{S}(\nu) \leftarrow \hat{S}_j(\nu)/(QM)$ , for all  $j$ 

```

result of the multilinear extension method and provide the bases of the two algorithms that we present in the next section.

IV. THE PROPOSED SAMPLING ALGORITHMS

In this section, we present two algorithms based on the multilinear extension method developed by Owen [7] (in his honor, we call the algorithms *Owen sampling*).

A. The Owen Sampling Algorithm

The first algorithm is based on a discretization of Equation (2). In this algorithm, the number of samples is controlled by two positive integers Q and M , and the Shapley values estimators are shown by $\hat{S}(\nu) = (\hat{S}_0(\nu), \dots, \hat{S}_n(\nu))$. Algorithm 1 explains the method in pseudo-code.

This Algorithm requires as input X, ν , and the number of samples which are controlled by the parameters Q and M . The parameter Q controls the level of discretization of the integral in Equation (2), and the parameter M controls the accuracy of the estimation of the expected value $e_j(q)$ in Equation (3). The optimal choices of Q and M are yet to be explored, however, of the two, the parameter Q plays a more significant role in the sense that by increasing Q , in general, we increase the total random numbers generated in the inner loop controlled by M . But the opposite of this does not occur. In all our experiments, we fix the value of $M = 2$, i.e., the total number of samples is equal to $2Q$.

The outer loop in Line 2 of this algorithm approximates the integral of Equation (2) by initializing zero values for its integrands e_j in Line 3. If we divide these values by M , i.e. e_j/M , they are approximations of $e_j(q)$ in Equation (3). More precisely, the inner loop in Line 4 uses a Monte Carlo estimation method that leads to the estimation of Equation (3). Since the probabilistic structure of the algorithm is governed by Bernoulli distribution, in Line 5, for each iteration of the inner loop, we generate $n + 1$ Bernoulli random numbers independently, which are temporarily stored in the vector $I_m^{(q)}$, i.e., each element of this vector is either zero or one depending on the outcome of the Bernoulli random number.

This vector is used to calculate the marginal contributions $h_{m,j}^{(q)}$. In Line 6, the \odot operator is the element-wise multiplication, and it is applied to the vectors $I_m^{(q)}$ and X . The operand $X^{(j)}$ in Line 6 indicates the vector $(0, 0, \dots, 0, x_j, 0, \dots, 0)$. The unnormalized approximated values of the integral in Equation (2) are then updated in Line 9. To normalize them we need to divided them by the total number of iterations QM . The convergence of this algorithm is guaranteed by the following proposition.

Proposition IV.1. *Suppose that for a fixed j , function $q \mapsto e_j(q)$ is Riemann (or Lebesgue) integrable¹, i.e., $\int_0^1 |e_j(q)| dq < \infty$. Then the estimators of Algorithm 1 converge to the true Shapley values as Q and M increase.*

Proof. See the Appendix. \square

B. Time Complexity Analysis

Algorithm 1 can be improved using parallel programming, however, an implementation method such as this is not our main criterion of comparison. Note that generating a Bernoulli random number is computable in constant time, i.e., $O(1)$ [18, 19]. Therefore, assuming that the time complexity of calculating $h_{m,j}^{(q)}$ is polynomial, the time complexity of the algorithm is polynomial. This has the same time complexity of the Castro sampling algorithm [1], but in order to have a fair comparison, a more detailed investigation is required.

In the Castro sampling algorithm, in each iteration three major operations are performed:

- 1) first, a random permutation from the set of n features is drawn using a uniform probability, the time-complexity of this is $O(n)$;
- 2) second for a given permutation, the set of predecessors of the features in that permutation is calculated. This is also $O(n)$, and;
- 3) finally, the value of the machine learning model is calculated.

By a one to one comparison of Castro sampling algorithm to Algorithm 1, in each iteration, we draw a random sample with time complexity $n \cdot O(1) = O(n)$, instead of calculating the predecessors of the features, we calculate the identity $e_j = e_j + h_{m,j}^{(q)}$, for all j , that is $O(n)$, and finally we calculate the value of the machine learning model. In Algorithm 1, a total of QM samplings is required, while in the Castro sampling algorithm, this is handled by a single parameter, M_c .

Therefore, for a fair comparison, for QM samplings of Algorithm 1, we have to consider $M_c = QM$ ones of the Castro sampling algorithm. However, as our experiments show, both the accuracy and the actual running time of Algorithm 1 are better than that of the Castro sampling algorithm. An important feature of Algorithm 1 is symmetry which can be exploited to improve the Algorithm 1 as follows.

¹This is a technical assumption, and it is satisfied in all practical circumstances.

Algorithm 2 Halved Owen Sampling

Input: $\nu, X = (x_0, x_1, \dots, x_n), Q, M$

Output: $\hat{S}(\nu) = (\hat{S}_0(\nu), \hat{S}_1(\nu), \dots, \hat{S}_n(\nu))$

- 1: $\hat{S}_j(\nu) := 0$, for all j
 - 2: **for** $q = 0, 1/Q, 2/Q, \dots, 0.5$ **do**
 - 3: $e_j := 0$, for all j
 - 4: **for** $m = 1, 2, \dots, M$ **do**
 - 5: $I_m^{(q)} \leftarrow (b_j : b_j \sim \text{Bern}(q), \text{ for all } j)$
 - 6: $I_m^{(-q)} \leftarrow 1 - I_m^{(q)}$
 - 7: $h_{m,j}^{(q)} \leftarrow \nu(I_m^{(q)} \odot X + X^{(j)}) - \nu(I_m^{(q)} \odot X)$, for all j
 - 8: $h_{m,j}^{(-q)} \leftarrow \nu(I_m^{(-q)} \odot X + X^{(j)}) - \nu(I_m^{(-q)} \odot X)$, for all j
 - 9: $e_j \leftarrow e_j + h_{m,j}^{(q)} + h_{m,j}^{(-q)}$, for all j
 - 10: **end for**
 - 11: $\hat{S}_j(\nu) \leftarrow \hat{S}_j(\nu) + e_j$, for all j
 - 12: **end for**
 - 13: $\hat{S}(\nu) \leftarrow \hat{S}(\nu)/(QM)$, for all j
-

C. Improving Owen Sampling through Symmetry

For $0 \leq q \leq 1$, we let $h_{m,j}^{(-q)}$ represent the calculation of the output of the machine learning model as in Line 6 of Algorithm 1, when instead of q , we use $1 - q$. Looking at Algorithm 1, it is clear that by using this symmetry, we only need to run the outer loop of this algorithm halfway through, i.e., for $q = 0, 1/Q, \dots, 0.5$, as for each q , we can duplicate the calculations of the output of the machine learning model by using $1 - q$. For instance, suppose that we only have three features (x_0, x_1, x_2) , and for a given q , we have generated the Bernoulli random sequence $(1, 1, 0)$ which means that the features x_0 and x_1 are selected according to this distribution, and feature x_2 is absent. This is equivalent to choosing only x_2 based on the random sequence $(0, 0, 1)$ generated using a Bernoulli distribution with parameter $1 - q$. This will have two important consequences, first, we only need to generate $MQ/2$ random numbers each with Bernoulli distribution and, second and most importantly, this will further reduce the variance of the estimator which will be discussed in our experimental section. For this reason, we call this algorithm *Halved Owen Sampling*. This is formalized in Algorithm 2.

Algorithm 2 is almost identical to Algorithm 1 except for few lines. As already discussed, comparing Line 2 of Algorithm 1 against Line 2 of Algorithm 2, we now only need to explore $Q/2$ iterations. In Lines 7 and 8, we calculate the output of the model for two cases, q and $1 - q$, using the vectors created in Lines 5 and 6. The result of Line 7 and 8 is then added to Line 9. The convergence of this algorithm is guaranteed by the following theorem.

Theorem IV.1. *Suppose that for a fixed j , function $q \mapsto e_j(q)$ is Riemann (or Lebesgue) integrable, i.e., $\int_0^1 |e_j(q)| dq < \infty$. Then the estimates of Algorithm 2 converge to the true Shapley values as Q and M increase.*

Proof. See the Appendix. \square

For a fair comparison of Algorithm 2 against the Castro sampling algorithm, we use similar arguments as done for Algorithm 1 where for a QM sample size of this algorithm, we consider an equivalent $M_c = QM$ sample size of the Castro sampling algorithm. This may sound counter-intuitive because it would appear that this new algorithm uses $QM/2$ samples, i.e., we generate $QMn/2$ random numbers in contrast to QMn random numbers of Algorithm 1. However, in its inner most loop we calculate the output of the machine learning model twice. Hence, the number of operations is still in the order of QM . Therefore we consider the same sample size as for Algorithm 1.

V. EXPERIMENTAL SETUP

In this section, we perform several experiments to test our algorithms. Although, the algorithms are applicable to any machine learning model, here we focus only on Multilayer Perceptrons (MLPs); further experiments on more complex networks are left to future work. The software used to run these experiments is available at the following link <https://github.com/aldolipani/OwenShap>.

A. Datasets and Preprocessing

We experiment with the following two datasets:

Credit Card Dataset (CC). This is a financial dataset which is a collection of credit card data used by Yeh and Lien [20]. There is a total number of 29,351 observations where each observation is made of 23 features and a binary target variable. The features are either financial (such as pay related information) or non-financial like age. The target variable is either zero or one with one indicating the default of the credit card account.

Modified NIST (MNIST). This is a large database of handwritten digits that is commonly used for training and testing machine learning models [21]. It was created by “re-mixing” the samples from another dataset. Each sample is a black and white image of a handwritten digit. Furthermore, the black and white images are normalized to fit into a 28x28 pixel bounding box. The MNIST dataset contains 70,000 images.

These two datasets have a different number of features (23 and 784). In order to compare the various sampling algorithms we need to compute the exact Shapely values for at least one dataset. However, both datasets make this calculation infeasible due to their feature sizes. For this reason, we reduce the dimensionality of the CC dataset by only selecting the first 15 features.

B. MLPs and Training Details

The activation function of the input layer and hidden layers is always a sigmoid, while the activation function of the output layer is a softmax. All MLPs are trained using Adam as optimizer with its default parameters and the binary cross-entropy as loss function. We use an early stopping criteria to avoid overfitting, where we stop training after the loss on the validation set has not improved over 3 epochs. We divide both datasets into 64% training set, 16% validation set and 20%

test set. The rest of the hyper-parameters, i.e., the number of hidden layers and number of neurons in each layer, are chosen via a Monte Carlo sampling of models using the training and validation sets.

For the CC dataset, we sample 1000 models with a number of hidden layers from 0 to 3 and a number of neurons for each hidden layer from 1 to 15. The best model has 2-hidden layers with 13 and 9 neurons each. The accuracy of this model on the CC test set is 0.8247. For the MNIST dataset, we sample 1000 models with a number of hidden layers from 0 to 3 and a number of neurons for each hidden layer from 25 to 500 at multiples of 25. The best model has 2-hidden layers with 300 and 25 neurons each. The accuracy of this model on the MNIST test set is 0.9818.

C. Experiments

Approximation Error. To evaluate the quality of the different sampling algorithms, we compare them against the exact Shapely values. Computing the Shapely values is an expensive operation because it grows exponentially with the number of features. In order to make this feasible, we only perform this analysis on the smaller CC dataset. Here, we take a sample of size 50 of the test set and compute the exact Shapely values. Then for each algorithm, as we increase the number of algorithms’ samplings, we compute the average of the Mean Squared Errors (MSEs) across the examples. In order to make a fair comparison, so that the numbers of samples of the various algorithms are equivalent, we set $M_c = 2m$ for the Castro sampling, where m is a positive integer, and $Q = m$ and $M = 2$ for the Owen sampling.

Variance Analysis. Besides the quality of the predictions, we perform a complementary analysis of the sampling algorithms by measuring their sample variances. To do this, we first take a sample of the test set of size 50. Then for each sampling algorithm and each example, we estimate the Shapely values increasing the sampling sizes from 2 to 200 at steps of 2. At each step, we compute the standard deviations of the feature estimates over the previous steps. Then, we compute the average across the features. Finally, we average across the examples. Note that the increase in the sample size is done by varying the equivalent number of samples, i.e., $M_c = 2Q$. As the number of samples increases, a decrease of the averaged sample standard deviations would confirm the convergence of the estimators. Note that from Proposition IV.1 and Theorem IV.1, the Owen sampling algorithms converge to the exact Shapely values only when both M and Q increase. By fixing $M = 2$, the algorithms converge to an estimation of the Shapely values.

VI. RESULTS AND DISCUSSION

In Figure 1, we observe the results of the error analysis. Both plots measure the MSE of the exact Shapely values against the outcome of the sampling algorithms. The box-plot is computed over 50 examples from the test set and using 2000 equivalent samples ($M_c = 2000$ for Castro, $Q = 1000$ and $M = 2$ for Owen). From this box-plot we can observe that the

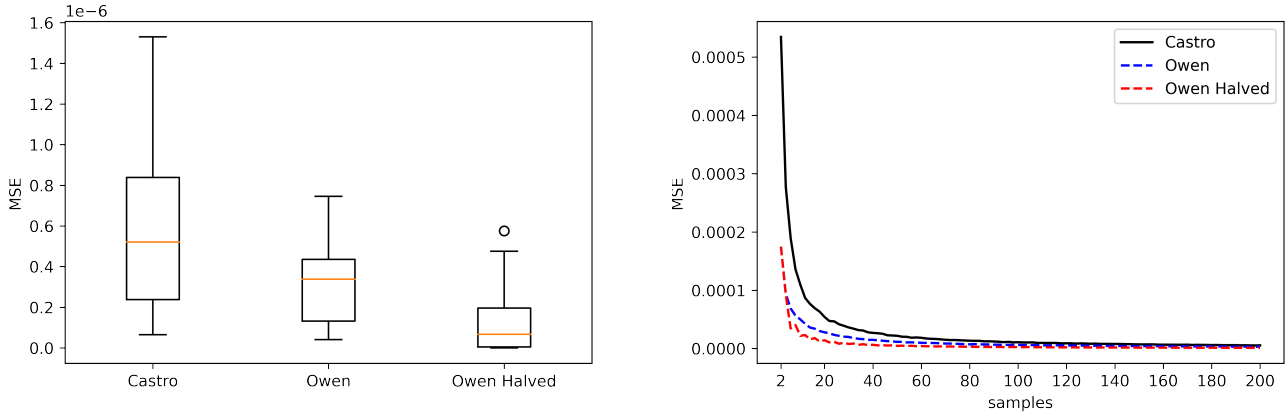


Fig. 1: MSE of the exact Shapely values against the Castro and Owen samplings. The box-plot shows the variance of the error over 50 examples and 2000 samples. The line-plot shows how the error decreases when the number of samples is increased.

Algorithm	Parameters	MSE (10^{-6})	Time (ms)
Castro	$M_c = 2000$	0.5575	3.004
Owen	$M = 2, Q = 1000$	0.3184	1.044
Halved Owen	$M = 2, Q = 1000$	0.1207	0.968

TABLE I: Summary of the results presented in box-plot in Fig. 1. Results are averaged over 50 examples of the CC test set.

Owen sampling outperforms the Castro sampling. Moreover, the Halved variant of the Owen sampling further improves the original Owen sampling. In Table I, we observe the summary of the results shown in the box-plot. Moreover, we also present the expected running-time per example of the algorithms. This result, although may depend on the CPU (AMD Threadripper 2950X CPU) and GPU (Nvidia Titan RTX) used, shows that the Owen algorithm achieved around a 3-fold speed-up against the Castro algorithm.

In Figure 1, we also graph a line-plot of the MSEs computed over 50 examples from the CC test set. The x-axis of this plot represents the number of equivalent samples. In this plot we observe that the MSE of the Owen sampling tends to zero faster than that of Castro sampling. This plot further confirms the superiority of the Owen sampling and in particular of its Halved version.

In Figure 2, we present two plots each computed on a different dataset, CC on the left and MNIST on the right. The x-axis and y-axis of these plots represent the number of equivalent samples and the sample standard deviation. As previously mentioned, since all the estimators converge, these sampled standard deviation tend to decrease. Here we can observe that as anticipated when discussing the previous experiment, the Castro sampling has a higher variance with respect to the Owen sampling. Moreover, we see that the Halved Owen sampling performs better than the Owen sampling. These results agree with what is observed in the previous experiment.

Finally, in Figure 3, we provide a qualitative analysis of the results of these sampling algorithms. This in order to observe whether the quality of these estimators are perceivable when producing saliency maps – a common way to show on what the model is basing its prediction. We do this for the MNIST dataset. In this case we compare the results of the Castro sampling against the Halved Owen sampling with an equivalent number of samples of 12. In order to have a reference, we use as a ground truth the results of the Castro sampling algorithm when computed with a very large number of samples, 5,000 in this case. These images’ pixels colored in blue represent the negative contributions and in red the positive ones towards the predicted score. When judging these images we should not focus on their background color because this depends on the maximum and minimum pixel values of the image. In these images we should focus instead on the most saturated pixels, these are in fact those that mostly contribute towards or against the model prediction. In these images we can observe that in most of the cases the Owen sampling identifies better those highly contributing pixels than the Castro sampling.

VII. CONCLUSION AND FUTURE WORK

We have provided a sampling algorithm to efficiently estimate Shapely values that can be also used as a ground truth for comparison purposes, since we have proved the convergence of this algorithm to the exact Shapely values. The method takes advantage of a variance reduction method and provide more accurate estimations for the Shapely values. In all the experiments that we have carried out, our comparison with the Castro algorithm [1] has been based on equal bases, i.e., we have used the equivalent number of samples. In our experiments, MLP models are used to fit with the datasets for classification problems (though regression problems could be considered as well). An optimal architecture of the MLPs are found using a cross validation analysis. Then we have




Fig. 2: Variance analysis of the estimators: the left-hand side plot is computed for the CC dataset, and the right-hand one for the MNIST dataset. Both plots are based on 50 examples of the test sets.




Fig. 3: Saliency maps of 10 randomly selected digits from the test set of the MNIST dataset. We use the top-most row as a reference to compare the results of the following two rows.

used these MLPs to test the effectiveness of our results on the corresponding datasets. First, we have started with a dataset with small number of features in which we can find the exact measurements of Shapley values. By selecting random instances from the test dataset and calculating the MSE in this case, our algorithms provide clear-cut improvements over the existing sampling algorithm leading to more accurate estimations. While an exact measurement of the Shapley values on datasets with a large number of features is not feasible, we have used an analysis of variance to show that our algorithm provide estimators for which their variance rapidly flattens leading to more accurate estimators. For consistency, the same analysis is also carried out for datasets with small number of features, and the same results are confirmed.

The experiments and the algorithms can be still improved in several directions which we leave to future work. More experimental analysis on different datasets could be carried out more complex deep learning architectures than MLPs, since our algorithm could work with any machine learning model. The accuracy of our algorithm is controlled by two parameters, in our analysis, we took one of the parameters to be equal to 2. However, more efficient and smart combinations of these parameters might improve the performance of our algorithms, in particular combining this with some statistical analysis, it might be possible to obtain the least number of operations required to reach certain accuracy through finding confidence intervals.

REFERENCES

- [1] J. Castro, D. Gómez, and J. Tejada, "Polynomial calculation of the shapley value based on sampling," *Computers & Operations Research*, vol. 36, no. 5, pp. 1726–1730, 2009.
- [2] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in neural information processing systems*, 2017, pp. 4765–4774.
- [3] M. T. Ribeiro, S. Singh, and C. Guestrin, "'why should i trust you?' explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.
- [4] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 3145–3153.
- [5] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, vol. 10, no. 7, p. e0130140, 2015.
- [6] L. S. Shapley, "A value for n-person games," *Contributions to the Theory of Games*, vol. 2, no. 28, pp. 307–317, 1953.
- [7] G. Owen, "Multilinear extensions of games," *Management Science*, vol. 18, no. 5-part-2, pp. 64–79, 1972.
- [8] S. Lipovetsky and M. Conklin, "Analysis of regression in game theory approach," *Applied Stochastic Models in Business and Industry*, vol. 17, no. 4, pp. 319–330, 2001.
- [9] S. S. Fatima, M. Wooldridge, and N. R. Jennings, "A linear approximation method for the shapley value," *Artificial Intelligence*, vol. 172, no. 14, pp. 1673–1699, 2008.
- [10] A. Datta, S. Sen, and Y. Zick, "Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems," in *2016 IEEE symposium on security and privacy (SP)*. IEEE, 2016, pp. 598–617.
- [11] M. Ancona, C. Oztireli, and M. Gross, "Explaining deep neural networks with a polynomial time algorithm for shapley value approximation," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 09–15 Jun 2019, pp. 272–281.
- [12] M. A. Jones and J. M. Wilson, "Multilinear extensions and values for multichoice games," *Mathematical Methods of Operations Research*, vol. 72, no. 1, pp. 145–169, 2010.
- [13] T. P. Michalak, K. V. Aadithya, P. L. Szczepanski, B. Ravindran, and N. R. Jennings, "Efficient computation of the shapley value for game-theoretic network centrality," *Journal of Artificial Intelligence Research*, vol. 46, pp. 607–650, 2013.
- [14] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2017, pp. 3319–3328.
- [15] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 3145–3153.
- [16] H. Moulin, *Fair division and collective welfare*. MIT press, 2004.
- [17] C. Izzo, A. Lipani, R. Okhrati, and F. Medda, "A baseline for Shapley values in MLPs: from missingness to neutrality," 2020.
- [18] K. Bringmann and T. Friedrich, "Exact and efficient generation of geometric random variates and random graphs," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2013, pp. 267–278.
- [19] V. Kachitvichyanukul and B. W. Schmeiser, "Binomial random variate generation," *Communications of the ACM*, vol. 31, no. 2, pp. 216–222, 1988.
- [20] I.-C. Yeh and C.-h. Lien, "The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients," *Expert Systems with Applications*, vol. 36, no. 2, pp. 2473–2480, 2009.
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

Appendix: A Multilinear Sampling Algorithm to Estimate Shapley Values

Ramin Okhrati
University College London
London, United Kingdom
r.okhrati@ucl.ac.uk

Aldo Lipani
University College London
London, United Kingdom
aldo.lipani@ucl.ac.uk

I. APPENDIX A

For the convenient of the reader, we review certain concepts and definitions of the paper. In what follows, ν represents the output of the machine learning model, $I_m^{(q)} \odot X$ is a vector obtained by the element-wise multiplications of the elements of $I_m^{(q)}$ and X , and $X^{(j)} := (0, 0, \dots, 0, x_j, 0, \dots, 0)$.

Proof of Proposition III.1. It is shown by [1] that the Shapley value of feature j is given by

$$S_j(v) = \int_0^1 e_j(q) dq, \quad (1)$$

where

$$e_j(q) = \mathbb{E}[\nu(E_j \cup \{j\}) - \nu(E_j)], \quad (2)$$

and E_j is the random subset of $X - \{x_j\}$ equipped with the aforementioned probabilistic structure in our work. Therefore, Equation (2) can be equivalently written as

$$e_j(q) = \mathbb{E}[\nu(\mathcal{I}^{(q)} \odot X + X^{(j)}) - \nu(\mathcal{I}^{(q)} \odot X)],$$

where $\mathcal{I}^{(q)}$ is a random vector with each element admits a Bernoulli distribution with parameter q .

Since the mapping $q \mapsto e_j(q)$ is Riemann integrable, it can be approximated using a Riemann sum as follows:

$$S_j(v) = \lim_{Q \rightarrow \infty} \sum_{i=0}^Q \frac{1}{Q} e_j(i/Q). \quad (3)$$

Furthermore, for a fixed value, i , $0 \leq i \leq Q$, $e_j(i/Q)$ can be approximated using Monte Carlo simulation, i.e.,

$$e_j(i/Q) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M (\nu(I_m^{(i/Q)} \odot X + X^{(j)}) - \nu(I_m^{(i/Q)} \odot X)),$$

where $I_m^{(i/Q)}$, for a fixed m , is a vector of Bernoulli random numbers with parameter i/Q .

Hence, Equation (3) is equal to

$$\frac{1}{Q \times M} \lim_{Q \rightarrow \infty} \lim_{M \rightarrow \infty} \sum_{i=0}^Q \sum_{m=1}^M (\nu(I_m^{(i/Q)} \odot X + X^{(j)}) - \nu(I_m^{(i/Q)} \odot X)). \quad (4)$$

Now, let us define the estimator $\hat{S}_j(\nu)$ by

$$\hat{S}_j(\nu) := \frac{1}{Q \times M} \sum_{i=0}^Q \sum_{m=1}^M (\nu(I_m^{(i/Q)} \odot X + X^{(j)}) - \nu(I_m^{(i/Q)} \odot X)).$$

This is the same estimator used in Algorithm 1 which converges to $S_j(\nu)$ by Equation (4). Therefore, we have proved that as Q and M increase $\hat{S}_j(\nu)$ converges to $S_j(\nu)$. \square

Proof of Theorem III.1. The proof follows from the same methodology as in the proof of Proposition III.1 which is based on Equation (1). This time, the idea is to write Equation (1) as follows:

$$\begin{aligned} S_j(v) &= \int_0^1 e_j(q) dq = \int_0^{0.5} e_j(q) dq + \int_{0.5}^1 e_j(q) dq \\ &= \int_0^{0.5} (e_j(q) + e_j(1-q)) dq. \end{aligned}$$

Now, similar to the proof of Proposition III.1, this integral can be approximated using the Riemann sum:

$$S_j(v) = \lim_{Q \rightarrow \infty} \sum_{i=0}^{Q/2} \frac{1}{Q} (e_j(i/n) + e_j(1-i/n)). \quad (5)$$

For a fixed value $0 \leq i \leq Q$, $e_j(i/n)$ and $e_j(1-i/n)$ can be approximated respectively using the Monte Carlo simulations

$$e_j(i/Q) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M (\nu(I_m^{(i/Q)} \odot X + X^{(j)}) - \nu(I_m^{(i/Q)} \odot X)),$$

and

$$e_j(1-i/Q) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M (\nu(I_m^{(1-i/Q)} \odot X + X^{(j)}) - \nu(I_m^{(1-i/Q)} \odot X)).$$

Therefore, Equation (5) can be written as

$$S_j(v) = \frac{1}{Q \times M} \lim_{Q \rightarrow \infty} \lim_{M \rightarrow \infty} \sum_{i=0}^{Q/2} \sum_{m=1}^M (h_{i/Q}^{j,m} + h_{1-i/Q}^{j,m}),$$

where

$$h_q^{j,m} = \nu(I_m^{(q)} \odot X + X^{(j)}) - \nu(I_m^{(q)} \odot X).$$

Finally, we define the estimator $\hat{S}_j(\nu)$ by

$$\hat{S}_j(\nu) = \frac{1}{Q \times M} \sum_{i=1}^{Q/2} \sum_{m=1}^M (h_{i/Q}^{j,m} + h_{1-i/Q}^{j,m}).$$

This is the same estimator used in Algorithm 2. So we have just proved that as Q and M increase, $\hat{S}_j(\nu)$ converges to $S_j(\nu)$.

□

REFERENCES

- [1] G. Owen, "Multilinear extensions of games," *Management Science*, vol. 18, no. 5-part-2, pp. 64–79, 1972.