Improving KernelSHAP: Practical Shapley Value Estimation via Linear Regression

Ian Covert
University of Washington

Su-In Lee
University of Washington

Abstract

The Shapley value concept from cooperative game theory has become a popular technique for interpreting ML models, but efficiently estimating these values remains challenging, particularly in the model-agnostic setting. Here, we revisit the idea of estimating Shapley values via linear regression to understand and improve upon this approach. By analyzing the original KernelSHAP alongside a newly proposed unbiased version, we develop techniques to detect its convergence and calculate uncertainty estimates. We also find that the original version incurs a negligible increase in bias in exchange for significantly lower variance, and we propose a variance reduction technique that further accelerates the convergence of both estimators. Finally, we develop a version of KernelSHAP for stochastic cooperative games that yields fast new estimators for two global explanation methods.

1 INTRODUCTION

Shapley values are central to many machine learning (ML) model explanation methods (e.g., SHAP, IME, QII, Shapley Effects, Shapley Net Effects, SAGE) [24, 25, 38, 13, 32, 23, 12]. Though developed in the cooperative game theory context [36], recent work shows that Shapley values provide a powerful tool for explaining how models work when either individual features [24], individual neurons in a neural network [18], or individual samples in a dataset [17] are viewed as players in a cooperative game. They have become a go-to solution for allocating credit and quantifying contributions due to to their appealing theoretical properties.

Proceedings of the 24[th] International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130.

The main challenge when using Shapley values is calculating them efficiently. A naive calculation has computational complexity that is exponential in the number of players, so numerous approaches have been proposed to accelerate their calculation. Besides brute-force methods [23], other techniques include sampling-based approximations [38, 37, 10, 12], model-specific approximations (e.g., TreeSHAP) [1, 25] and a linear regression-based approximation (KernelSHAP) [24].

Here, we revisit the regression-based approach to address several shortcomings in KernelSHAP. Recent work has questioned whether KernelSHAP is an unbiased estimator [29], and, unlike sampling-based estimators [6, 26, 12], KernelSHAP does not provide uncertainty estimates. Furthermore, it provides no guidance on the number of samples required because its convergence properties are not well understood.

We address each of these problems, in part by building on a newly proposed unbiased version of the regression-based approach. Our contributions include:

1. Deriving an unbiased version of KernelSHAP and showing empirically that the original version incurs a negligible increase in bias in exchange for significantly lower variance
2. Showing how to detect KernelSHAP's convergence, automatically determine the number of samples required, and calculate uncertainty estimates for the results
3. Proposing a variance reduction technique that further accelerates KernelSHAP's convergence
4. Adapting the regression-based approach to stochastic cooperative games [9] to provide fast new approximations for two global explanation methods, SAGE [12] and Shapley Effects [32]

With these new insights and tools, we offer a more practical approach to Shapley value estimation via linear regression.[1]

[1] `https://github.com/iancovert/shapley-regression`

2 THE SHAPLEY VALUE

We now provide background information on cooperative game theory and the Shapley value.

2.1 Cooperative Games

A *cooperative game* is a function $v : 2^d \mapsto \mathbb{R}$ that returns a value for each coalition (subset) $S \subseteq D$, where $D = \{1, \dots, d\}$ represents a set of *players*. Cooperative game theory has become increasingly important in ML because many methods frame model explanation problems in terms of cooperative games [11]. Notably, SHAP [24], IME [38] and QII [13] define cooperative games that represent an individual prediction's dependence on different features. For a model f and an input x, SHAP (when using the marginal distribution [24]) analyzes the cooperative game v_x, defined as

$$v_x(S) = \mathbb{E}[f(x_S, X_{D\setminus S})], \tag{1}$$

where $x_S \equiv \{x_i : i \in S\}$ represents a feature subset and X_S is the corresponding random variable. Two other methods, Shapley Effects [32] and SAGE [12], define cooperative games that represent a model's behavior across the entire dataset. For example, given a loss function ℓ and response variable Y, SAGE uses a cooperative game w that represents the model's predictive performance given a subset of features X_S:

$$w(S) = -\mathbb{E}\Big[\ell\big(\mathbb{E}[f(X) \mid X_S], Y\big)\Big]. \tag{2}$$

Several other techniques also frame model explanation questions in terms of cooperative games, where a target quantity (e.g., model loss) varies as groups of players (e.g., features) are removed, and the Shapley value summarizes each player's contribution [11].

2.2 Shapley Values

The Shapley value [36] assumes that the *grand coalition* D is participating and seeks to provide each player with a fair allocation of the total profit, which is represented by $v(D)$. Fair allocations must be based on each player's contribution to the profit, but a player's contribution is often difficult to define. Player i's *marginal contribution* to the coalition S is the difference $v(S \cup \{i\}) - v(S)$, but the marginal contribution typically depends on which players S are already participating.

The Shapley value resolves this problem by deriving a unique value based on a set of fairness axioms; see [36, 30] for further detail. It can be understood as a player's average marginal contribution across all possible player orderings, and each player's Shapley value $\phi_1(v), \dots, \phi_d(v)$ for a game v is given by:

$$\phi_i(v) = \frac{1}{d} \sum_{S \subseteq D\setminus\{i\}} \binom{d-1}{|S|}^{-1} \Big(v(S\cup\{i\}) - v(S)\Big). \tag{3}$$

Many ML model explanation methods can be understood in terms of ideas from cooperative game theory [11], but the Shapley value is especially popular and is also widely used in other fields [2, 33, 39].

2.3 Weighted Least Squares Characterization

While we can characterize the Shapley value in many ways, the perspective most relevant to our work is viewing it as a solution to a weighted least squares problem. Many works have considered fitting simple models to cooperative games [8, 20, 19, 14, 15, 28], particularly additive models of the form

$$u(S) = \beta_0 + \sum_{i \in S} \beta_i.$$

Such additive models are known as *inessential games*, and although a game v may not be inessential, an inessential approximation can help summarize each player's average contribution. Several works [8, 20, 14] model games by solving a weighted least squares problem using a weighting function μ:

$$\min_{\beta_0,\dots,\beta_d} \sum_{S \subseteq D} \mu(S)\Big(u(S) - v(S)\Big)^2.$$

Perhaps surprisingly, different weighting kernels μ lead to recognizable optimal regression coefficients $(\beta_1^*, \dots \beta_d^*)$ [11]. In particular, a carefully chosen weighting kernel yields optimal regression coefficients equal to the Shapley values [8, 24]. The Shapley kernel μ_{Sh} is given by

$$\mu_{\text{Sh}}(S) = \frac{d-1}{\binom{d}{|S|}|S|(d-|S|)},$$

where the values $\mu_{\text{Sh}}(\{\}) = \mu_{\text{Sh}}(D) = \infty$ effectively enforce constraints $\beta_0 = v(\{\})$ for the intercept and $\sum_{i\in D} \beta_i = v(D) - v(\{\})$ for the sum of the coefficients. Lundberg and Lee [24] used this Shapley value interpretation when developing an approach to approximate SHAP values via linear regression.

3 LINEAR REGRESSION APPROXIMATIONS

As noted, Shapley values are difficult to calculate because they require examining each player's marginal contribution to every possible subset (Eq. 3). This leads to run-times that are exponential in the number of players, so efficient approximations are of great practical importance [38, 37, 24, 10, 1, 25, 12]. Here, we revisit the regression-based approach presented by Lundberg and Lee (KernelSHAP) [24] and then present an unbiased version of this approach whose properties are simpler to analyze.

3.1 Optimization Objective

The least squares characterization of the Shapley value suggests that we can calculate the values $\phi_1(v), \ldots, \phi_d(v)$ by solving the optimization problem

$$\begin{aligned} \min_{\beta_0,\ldots,\beta_d} \; & \sum_{0<|S|<d} \mu_{\text{Sh}}(S)\Big(\beta_0 + \sum_{i\in S}\beta_i - v(S)\Big)^2 \\ \text{s.t.} \quad & \beta_0 = v(\{\}), \quad \beta_0 + \sum_{i=1}^{d}\beta_i = v(D). \end{aligned} \tag{4}$$

Notation. We introduce new notation to make the problem easier to solve. First, we denote the non-intercept coefficients as $\beta = (\beta_1, \ldots, \beta_d) \in \mathbb{R}^d$. Next, we denote each subset $S \subseteq D$ using the corresponding binary vector $z \in \{0,1\}^d$, and with tolerable abuse of notation we write $v(z) \equiv v(S)$ and $\mu_{\text{Sh}}(z) \equiv \mu_{\text{Sh}}(S)$ for $S = \{i : z_i = 1\}$. Lastly, we denote a distribution over Z using $p(z)$, where we define $p(z) \propto \mu_{\text{Sh}}(z)$ when $0 < \mathbf{1}^T z < d$ and $p(z) = 0$ otherwise. With this, we can rewrite the optimization problem as

$$\begin{aligned} \min_{\beta_0,\ldots,\beta_d} \; & \sum_{z} p(z)\Big(v(\mathbf{0}) + z^T\beta - v(z)\Big)^2 \\ \text{s.t.} \quad & \mathbf{1}^T\beta = v(\mathbf{1}) - v(\mathbf{0}). \end{aligned} \tag{5}$$

3.2 Dataset Sampling

Solving the problem in Eq. 5 requires evaluating the cooperative game v with all 2^d coalitions. Evaluating $v(\mathbf{0})$ and $v(\mathbf{1})$ is sufficient to ensure that the constraints are satisfied, but all values $v(z)$ for z such that $0 < \mathbf{1}^T z < d$ are required to fit the model exactly. KernelSHAP manages this challenge by subsampling a dataset and optimizing an approximate objective. We refer to this approach as *dataset sampling*. Using n independent samples $z_i \sim p(Z)$ and their values $v(z_i)$, KernelSHAP solves the following problem:

$$\begin{aligned} \min_{\beta_0,\ldots,\beta_d} \; & \frac{1}{n}\sum_{i=1}^{n}\Big(v(\mathbf{0}) + z_i^T\beta - v(z_i)\Big)^2 \\ \text{s.t.} \quad & \mathbf{1}^T\beta = v(\mathbf{1}) - v(\mathbf{0}). \end{aligned} \tag{6}$$

The dataset sampling approach, also applied by LIME [35], offers the flexibility to use only enough samples to accurately approximate the objective. Given a set of samples $(z_1, \ldots, z_n)$, solving this problem is straightforward. The Lagrangian with multiplier $\nu \in \mathbb{R}$ is given by:

$$\begin{aligned} \hat{\mathcal{L}}(\beta,\nu) = \; & \beta^T\Big(\frac{1}{n}\sum_{i=1}^{n} z_i z_i^T\Big)\beta \\ & - 2\beta^T\Big(\frac{1}{n}\sum_{i=1}^{n} z_i\big(v(z_i) - v(\mathbf{0})\big)\Big) \\ & + \frac{1}{n}\sum_{i=1}^{n}\big(v(z_i) - v(\mathbf{0})\big)^2 \\ & + 2\nu\big(\mathbf{1}^T\beta - v(\mathbf{1}) + v(\mathbf{0})\big). \end{aligned}$$

If we introduce the shorthand notation

$$\hat{A}_n = \frac{1}{n}\sum_{i=1}^{n} z_i z_i^T \quad \text{and} \quad \hat{b}_n = \frac{1}{n}\sum_{i=1}^{n} z_i\Big(v(z_i) - v(\mathbf{0})\Big),$$

then we can use the problem's KKT conditions [5] to derive the following solution:

$$\hat{\beta}_n = \hat{A}_n^{-1}\Big(\hat{b}_n - \mathbf{1}\frac{\mathbf{1}^T\hat{A}_n^{-1}\hat{b}_n - v(\mathbf{1}) + v(\mathbf{0})}{\mathbf{1}^T\hat{A}_n^{-1}\mathbf{1}}\Big). \tag{7}$$

This method is known as KernelSHAP [24], and the implementation in the SHAP repository[2] also allows for regularization terms in the approximate objective (Eq. 6), such as the ℓ_1 penalty [40]. While this approach is intuitive and simple to implement, the estimator $\hat{\beta}_n$ is surprisingly difficult to characterize. As we show in Section 4, it is unclear whether it is unbiased, and understanding its variance and rate of convergence is not straightforward. Therefore, we derive an alternative approach that is simpler to analyze.

3.3 An Exact Estimator

Consider the solution to the problem that uses all 2^d player coalitions (Eq. 5). Rather than finding an *exact solution to an approximate problem* (Section 3.2),

[2] `http://github.com/slundberg/shap`

we now derive an *approximate solution to the exact problem*. The full problem's Lagrangian is given by

$$\begin{aligned}\mathcal{L}(\beta,\nu) = \beta^T\mathbb{E}[ZZ^T]\beta \\ - 2\beta^T\mathbb{E}\Big[Z\big(v(Z)-v(\mathbf{0})\big)\Big] \\ + \mathbb{E}\Big[\big(v(Z)-v(\mathbf{0})\big)^2\Big] \\ + 2\nu\big(\mathbf{1}^T\beta - v(\mathbf{1}) + v(\mathbf{0})\big),\end{aligned}$$

where we now consider Z to be a random variable distributed according to $p(Z)$. Using the shorthand notation

$$A = \mathbb{E}[ZZ^T] \quad \text{and} \quad b = \mathbb{E}\Big[Z\big(v(Z)-v(\mathbf{0})\big)\Big],$$

we can write the solution to the exact problem as:

$$\beta^* = A^{-1}\Big(b - \mathbf{1}\frac{\mathbf{1}^TA^{-1}b - v(\mathbf{1}) + v(\mathbf{0})}{\mathbf{1}^TA^{-1}\mathbf{1}}\Big). \tag{8}$$

Due to our setup of the optimization problem, we have the property that $\beta_i^* = \phi_i(v)$. Unfortunately, we cannot evaluate this expression in practice without evaluating v for all 2^d coalitions $S \subseteq D$.

However, knowledge of $p(Z)$ means that $A \in \mathbb{R}^{d\times d}$ can be calculated exactly and efficiently. To see this, note that $(ZZ^T)_{ij} = Z_iZ_j = \mathbb{1}(Z_i = Z_j = 1)$. Therefore, to calculate A, we need to estimate only $p(Z_i = 1)$ for diagonal values A_{ii} and $p(Z_i = Z_j = 1)$ for off-diagonal values A_{ij}. See Appendix A for their derivations.

Since b cannot be calculated exactly and efficiently due to its dependence on v, this suggests that we should use A's exact form and approximate β^* by estimating (only) b. We propose the following estimator for b:

$$\bar{b}_n = \frac{1}{n}\sum_{i=1}^{n} z_iv(z_i) - \mathbb{E}[Z]v(\mathbf{0}).$$

Using this, we arrive at an alternative to the original KernelSHAP estimator, which we refer to as *unbiased KernelSHAP*:

$$\bar{\beta}_n = A^{-1}\Big(\bar{b}_n - \mathbf{1}\frac{\mathbf{1}^TA^{-1}\bar{b}_n - v(\mathbf{1}) + v(\mathbf{0})}{\mathbf{1}^TA^{-1}\mathbf{1}}\Big). \tag{9}$$

In the next section, we compare these two approaches both theoretically and empirically.

4 ESTIMATOR PROPERTIES

We now analyze the consistency, bias and variance properties of the Shapley value estimators, and we consider how to detect, forecast, and accelerate their convergence.

4.1 Consistency, Bias and Variance

A *consistent* estimator is one that converges to the correct Shapley values β^* given a sufficiently large number of samples. If the game v has bounded value, then the strong law of large numbers implies that

$$\lim_{n\to\infty}\hat{A}_n = A \quad \text{and} \quad \lim_{n\to\infty}\hat{b}_n = \lim_{n\to\infty}\bar{b}_n = b,$$

where the convergence is almost sure. From this, we see that both estimators are consistent:

$$\lim_{n\to\infty}\hat{\beta}_n = \lim_{n\to\infty}\bar{\beta}_n = \beta^*.$$

Next, an *unbiased* estimator is one whose expectation is equal to the correct Shapley values β^*. This is difficult to verify for the KernelSHAP estimator $\hat{\beta}_n$ due to the interaction between $\hat{A}_n$ and $\hat{b}_n$ (see Eq. 7). Both $\hat{A}_n$ and $\hat{b}_n$ are unbiased, but terms such as $\mathbb{E}[\hat{A}_n^{-1}\hat{b}_n]$ and $\mathbb{E}[\hat{A}_n^{-1}\mathbf{1}\mathbf{1}^T\hat{A}_n^{-1}\hat{b}_n/(\mathbf{1}^T\hat{A}_n^{-1}\mathbf{1})]$ are difficult to characterize. To make any claims about KernelSHAP's bias, we rely instead on empirical observations.

In contrast, it is easy to see that the alternative estimator $\bar{\beta}_n$ is unbiased. Because of its linear dependence on $\bar{b}_n$ and the fact that $\mathbb{E}[\bar{b}_n] = b$, we can see that

$$\mathbb{E}[\bar{\beta}_n] = \beta^*.$$

We therefore conclude that the alternative estimator $\bar{\beta}_n$ is both consistent and unbiased, whereas the original KernelSHAP ($\hat{\beta}_n$) is only provably consistent. It is for this reason that we refer to $\bar{\beta}_n$ as *unbiased KernelSHAP*.

Regarding the estimators' variance, unbiased KernelSHAP is once again simpler to characterize. The values $\bar{\beta}_n$ are a function of $\bar{b}_n$, and the multivariate central limit theorem (CLT) [41] asserts that $\bar{b}_n$ converges in distribution to a multivariate Gaussian, or

$$\bar{b}_n\sqrt{n} \xrightarrow{D} \mathcal{N}(b, \Sigma_{\bar{b}}), \tag{10}$$

where $\Sigma_{\bar{b}} = \mathrm{Cov}\big(Zv(Z)\big)$. This implies that for the estimator $\bar{\beta}_n$, we have the convergence property

$$\bar{\beta}_n\sqrt{n} \xrightarrow{D} \mathcal{N}(\beta^*, \Sigma_{\bar{\beta}}), \tag{11}$$

where, due to its linear dependence on $\bar{b}_n$ (see Eq. 9), we have the covariance $\Sigma_{\bar{\beta}}$ given by

$$\Sigma_{\bar{\beta}} = C\Sigma_{\bar{b}}C^T \tag{12}$$

$$C = A^{-1} - \frac{A^{-1}\mathbf{1}\mathbf{1}^T A^{-1}}{\mathbf{1}^T A^{-1}\mathbf{1}}. \tag{13}$$

This allows us to reason about unbiased KernelSHAP's asymptotic distribution. In particular, we remark that $\bar{\beta}_n$ has variance that reduces at a rate of $\mathcal{O}(\frac{1}{n})$.

In comparison, the original KernelSHAP estimator $\hat{\beta}_n$ is difficult to analyze due to the interaction between the $\hat{A}_n$ and $\hat{b}_n$ terms. We can apply the CLT to either term individually, but reasoning about $\hat{\beta}_n$'s distribution or variance remains challenging.

To facilitate our analysis of KernelSHAP, we present a simple experiment to compare the two estimators. We approximated the SHAP values for an individual prediction in the census income dataset [22] and empirically calculated the mean squared error relative to the true SHAP values[3] across 250 runs. We then decomposed the error into bias and variance terms as follows:

$$\underbrace{\mathbb{E}\big[||\hat{\beta}_n - \beta^*||^2\big]}_{\text{Error}} = \underbrace{\mathbb{E}\big[||\hat{\beta}_n - \mathbb{E}[\hat{\beta}_n]||^2\big]}_{\text{Variance}} + \underbrace{\big|\big|\mathbb{E}[\hat{\beta}_n] - \beta^*\big|\big|^2}_{\text{Bias}}.$$

Figure 1 shows that the error for both estimators is dominated by variance rather than bias.[4] It also shows that KernelSHAP incurs virtually no bias in exchange for significantly lower variance. In Appendix F, we provide global measures of the bias and variance to confirm these observations across multiple examples and two other datasets. This suggests that although KernelSHAP is more difficult to analyze theoretically, it should be used in practice because its bias is negligible and it converges faster.

4.2 Variance Reduction via Paired Sampling

Having analyzed each estimator's properties, we now consider whether their convergence can be accelerated. We propose a simple variance reduction technique that leads to significantly faster convergence in practice.

When sampling n subsets according to the distribution $z_i \sim p(Z)$, we suggest a *paired sampling* strategy where each sample z_i is paired with its complement[5] $\mathbf{1} - z_i$. To show why this approach accelerates convergence, we focus on unbiased KernelSHAP, which is easier to analyze theoretically.

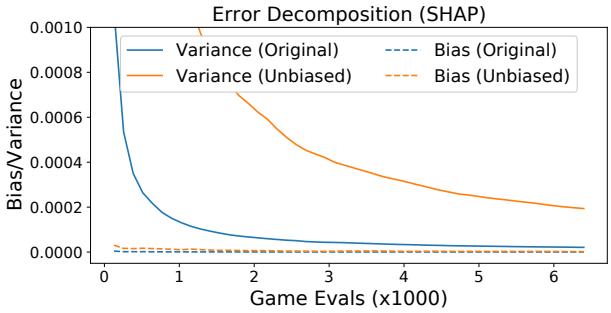

Figure 1: SHAP error decomposition for the original and unbiased KernelSHAP estimators. The bias and variance are calculated empirically across 250 runs.

When estimating b for unbiased KernelSHAP ($\bar{\beta}_n$), consider using the following modified estimator that combines z_i with $\mathbf{1} - z_i$:

$$\check{b}_n = \frac{1}{2n}\sum_{i=1}^{n}\big(z_i v(z_i) + (\mathbf{1}-z_i)v(\mathbf{1}-z_i) - v(\mathbf{0})\big). \tag{14}$$

Substituting this into unbiased KernelSHAP (Eq. 9) yields a new estimator $\check{\beta}_n$ that preserves the properties of being both consistent and unbiased:

$$\check{\beta}_n = A^{-1}\Big(\check{b}_n - \mathbf{1}\frac{\mathbf{1}^T A^{-1}\check{b}_n - v(\mathbf{1}) + v(\mathbf{0})}{\mathbf{1}^T A^{-1}\mathbf{1}}\Big). \tag{15}$$

For games v that satisfy a specific condition, we can guarantee that this sampling approach leads to $\check{\beta}_n$ having lower variance than $\bar{\beta}_n$, even when we account for $\check{b}_n$ requiring twice as many cooperative game evaluations as $\bar{b}_n$ (see proof in Appendix B).

Theorem 1. *The difference between the covariance matrices for the estimators $\bar{\beta}_{2n}$ and $\check{\beta}_n$ is given by*

$$\mathrm{Cov}(\bar{\beta}_{2n}) - \mathrm{Cov}(\check{\beta}_n) = \frac{1}{2n}CG_v C^T,$$

where G_v is a property of the game v, defined as

$$G_v = -\mathrm{Cov}\Big(Zv(Z), (\mathbf{1}-Z)v(\mathbf{1}-Z)\Big).$$

For sufficiently large n, $G_v \succeq 0$ guarantees that the Gaussian confidence ellipsoid $\bar{E}_{2n,\alpha}$ for $\bar{\beta}_{2n}$ contains the corresponding confidence ellipsoid $\check{E}_{n,\alpha}$ for $\check{\beta}_n$, or $\check{E}_{n,\alpha} \subseteq \bar{E}_{2n,\alpha}$, at any confidence level $\alpha \in (0,1)$.

[3] The true SHAP values use a sufficient number of samples to ensure convergence (see Section 4.3).

[4] The unbiased approach appears to have higher bias due to estimation error, but its bias is provably zero.

[5] We call $\mathbf{1} - z$ the *complement* because it is the binary vector for $D \setminus S$, where S corresponds to z.

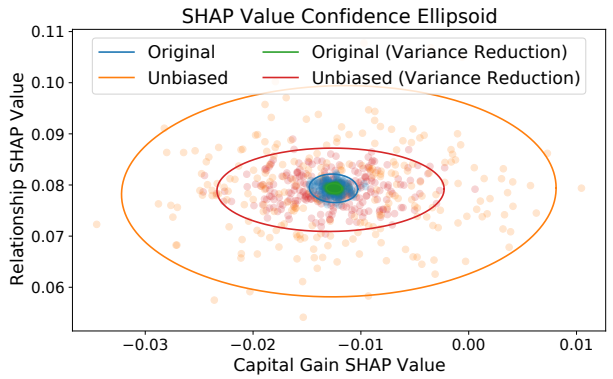

Figure 2: Gaussian 95% confidence ellipsoids for two SHAP values in a census income prediction (from 250 runs). The estimators use an equal number of samples.

Theorem 1 shows that $\check{\beta}_n$ is a more precise estimator than $\bar{\beta}_{2n}$ when the condition $G_v \succeq 0$ is satisfied (i.e., G_v is positive semi-definite). This may not hold in the general case, but in Appendix B we show that a weaker condition holds for *all games*: the diagonal values of G_v satisfy $(G_v)_{ii} \geq 0$ for any game v. Geometrically, this weaker condition means that $\bar{E}_{2n,\alpha}$ extends beyond $\check{E}_{n,\alpha}$ in the axis-aligned directions.

Figure 2 illustrates the result of Theorem 1 by showing empirical 95% confidence ellipsoids for two SHAP values. Although a comparable condition is difficult to derive for the original KernelSHAP estimator ($\hat{\beta}_n$), we find that the paired sampling approach yields a similar reduction in variance. Our experiments provide further evidence that this approach accelerates convergence for both estimators (Section 6).

4.3 Convergence Detection and Forecasting

One of KernelSHAP's practical shortcomings is its lack of guidance on the number of samples required to obtain accurate estimates. We address this problem by developing an approach for convergence detection and forecasting.

Previously, we showed that unbiased KernelSHAP ($\bar{\beta}_n$) has variance that reduces at a rate $\mathcal{O}(\frac{1}{n})$ (Eq. 10). Furthermore, its variance is simple to estimate in practice: we require only an empirical estimate $\hat{\Sigma}_{\bar{b}}$ of $\Sigma_{\bar{b}}$ (defined above), which we can calculate using an online algorithm, such as Welford's [42].

We also showed that the original KernelSHAP ($\hat{\beta}_n$) is difficult to characterize, but its variance is empirically lower than the unbiased version. Understanding its variance is useful for convergence detection, so we propose an approach for approximating it. Based on the results in Figure 1, we may hypothesize that KernelSHAP's variance reduces at the same rate of $\mathcal{O}(\frac{1}{n})$; in Appendix F, we examine this by plotting the product of the variance and the number of samples over the course of estimation. We find that the product is *constant* as the sample number increases, which suggests that the $\mathcal{O}(\frac{1}{n})$ rate holds in practice. This property is difficult to prove formally, but it can be used for simple variance approximation.

When running KernelSHAP, we suggest estimating the variance by selecting an intermediate value m such that $m << n$ and calculating multiple independent estimates $\hat{\beta}_m$ while accumulating samples for $\hat{\beta}_n$. For any n, we can then approximate $\mathrm{Cov}(\hat{\beta}_n)$ as

$$\mathrm{Cov}(\hat{\beta}_n) \approx \frac{m}{n}\mathrm{Cov}(\hat{\beta}_m),$$

where $\mathrm{Cov}(\hat{\beta}_m)$ is estimated empirically using the multiple independent estimates $\hat{\beta}_m$. This online approach has a negligible impact on the algorithm's run-time, and the covariance estimate can be used to provide confidence intervals for the final results.

Whether we use the original or unbiased version of KernelSHAP, the estimator's covariance at a given value of n lets us both detect and forecast convergence. For detection, we propose stopping at the current value n when the largest standard deviation is a sufficiently small portion t (e.g., $t = 0.01$) of the gap between the largest and smallest Shapley value estimates. For unbiased KernelSHAP, this criterion is equivalent to:

$$\max_i \sqrt{\frac{1}{n}(\hat{\Sigma}_{\bar{\beta}})_{ii}} < t\Big(\max_i\,(\bar{\beta}_n)_i - \min_i\,(\bar{\beta}_n)_i\Big).$$

To forecast the number of samples required to reach convergence, we again invoke the property that estimates have variance that reduces at a rate of $\mathcal{O}(\frac{1}{n})$. Given a value t and estimates $\bar{\beta}_n$ and $\hat{\Sigma}_{\bar{\beta}}$, the approximate number of samples $\hat{N}$ required is:

$$\hat{N} = \frac{1}{t^2}\Big(\frac{\max_i \sqrt{(\hat{\Sigma}_{\bar{\beta}})_{ii}}}{\max_i\,(\bar{\beta}_n)_i - \min_i\,(\bar{\beta}_n)_i}\Big)^2.$$

This allows us to forecast the time until convergence at any point during the algorithm. The forecast is expected to become more accurate as the estimated terms become more precise.

Our approach is theoretically grounded for unbiased KernelSHAP, and the approximate approach for the standard version of KernelSHAP relies only on the assumption that $\mathrm{Cov}(\hat{\beta}_n)$ reduces at a rate of $\mathcal{O}(\frac{1}{n})$. Appendix G shows algorithms for both approaches, which illustrate both variance reduction and convergence detection techniques.

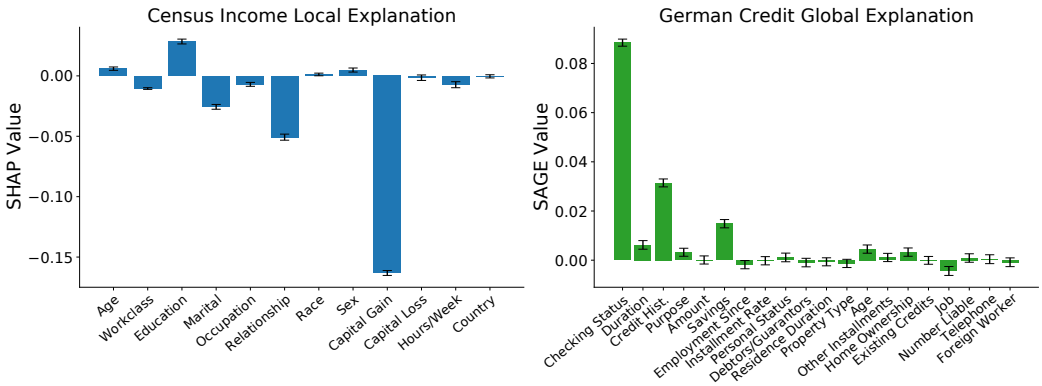

Figure 3: Shapley value-based explanations with 95% uncertainty estimates. Left: SHAP values for a single prediction with the census income dataset. Right: SAGE values for the German credit dataset.

5 STOCHASTIC COOPERATIVE GAMES

We have thus far focused on developing a regression-based approach to estimate Shapley values for any cooperative game. We now discuss how to adapt this approach to stochastic cooperative games, which leads to fast estimators for two global explanation methods.

5.1 Stochastic Cooperative Games

Stochastic cooperative games return a random value for each coalition of participating players $S \subseteq D$. Such games are represented by a function V that maps coalitions to a *distribution* of possible outcomes, so that $V(S)$ is a random variable [9, 7].

To aid our presentation, we assume that the uncertainty in the game can be represented by an exogenous random variable U. The game can then be denoted by $V(S, U)$, where $V(\cdot, U)$ is a deterministic function of S for any fixed value of the variable U.

Stochastic cooperative games provide a useful tool for understanding two global explanation methods, SAGE [12] and Shapley Effects [32]. To see why, assume an exogenous variable $U = (X, Y)$ that represents a random input-label pair, and consider the following game:

$$W(S, X, Y) = -\ell\Big(\mathbb{E}\big[f(X)|X_S\big], Y\Big). \tag{16}$$

The game W evaluates the (negated) loss with respect to the label Y given a prediction that depends only on the features X_S. The cooperative game used by SAGE can be understood as the expectation of this game, or $w(S) = \mathbb{E}_{XY}\big[W(S, X, Y)\big]$ (see Eq. 2). Shapley Effects is based on the expectation of a similar game, where the loss is evaluated with respect to the full model prediction $f(X)$ (see Appendix C). As we show next, an approximation approach tailored to this setting yields significantly faster estimators for these methods.

5.2 Generalizing the Shapley Value

It is natural to assign values to players in stochastic cooperative games like we do for deterministic games. We propose a simple generalization of the Shapley value for games $V(S, U)$ that averages a player's marginal contributions over both (i) player orderings and (ii) values of the exogenous variable U:

$$\phi_i(V) = \frac{1}{d} \sum_{S \subseteq D \setminus \{i\}} \binom{d-1}{|S|}^{-1} \mathbb{E}_U\Big[V(S \cup \{i\}, U) - V(S, U)\Big].$$

Due to the linearity property of Shapley values [36, 30], the following sets of values are equivalent:

1. The Shapley values of the game's expectation $\bar{v}(S) = \mathbb{E}_U[V(S, U)]$, or $\phi_i(\bar{v})$
2. The expected Shapley values of games with fixed U, or $\mathbb{E}_U[\phi_i(v_U)]$ where $v_u(S) = V(S, u)$
3. Our generalization of Shapley values to the stochastic cooperative game $V(S, U)$, or $\phi_i(V)$

The first two list items suggest ways of calculating the values $\phi_i(V)$ using tools designed for deterministic cooperative games. However, the expectation $\mathbb{E}_U\big[V(S, U)\big]$ may be slow to evaluate (e.g., if it is across an entire dataset), and calculating Shapley values separately for each value of U would be intractable if U has many possible values. We therefore introduce a third approach.

5.3 Shapley Value Approximation for Stochastic Cooperative Games

We now propose a fast, regression-based approach for calculating the generalized Shapley values $\phi_i(V)$ of

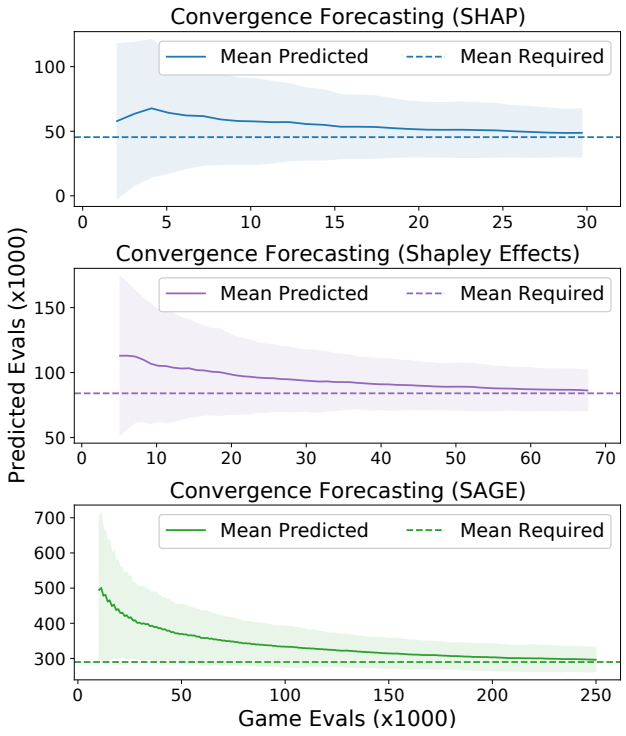

Figure 4: Convergence forecasting for SHAP, Shapley Effects and SAGE. The required number of samples is compared with the predicted number across 100 runs (with 90% confidence intervals displayed).

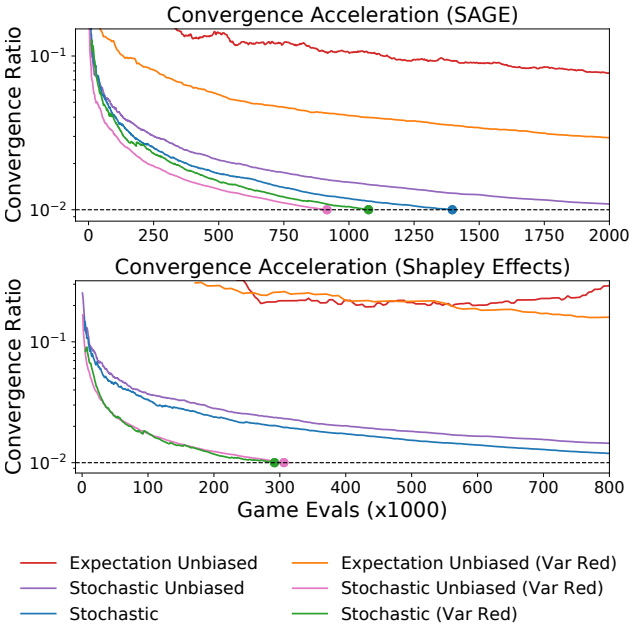

Figure 5: Convergence acceleration for SAGE and Shapley Effects. The ratio of the maximum standard deviation to the gap between the largest and smallest Shapley values is compared across six estimators.

stochastic cooperative games $V(S, U)$. Fortunately, it requires only a simple modification of the preceding approaches.

First, we must calculate the values $\mathbb{E}_U\big[V(\mathbf{1}, U)\big]$ and $\mathbb{E}_U\big[V(\mathbf{0}, U)\big]$ for the grand coalition and the empty coalition. Next, we replace our previous b estimators ($\hat{b}_n$ and $\bar{b}_n$) with estimators that use n pairs of independent samples $z_i \sim p(Z)$ and $u_i \sim p(U)$. To adapt the original KernelSHAP to this setting, we use

$$\tilde{b}_n = \frac{1}{2}\sum_{i=1}^{n} z_i\big(V(z_i, u_i) - \mathbb{E}_U\big[V(\mathbf{0}, U)\big]\big).$$

We then substitute this into the KernelSHAP estimator, as follows:

$$\tilde{\beta}_n = \hat{A}_n^{-1}\Big(\tilde{b}_n - \mathbf{1}\frac{\mathbf{1}^T\hat{A}_n^{-1}\tilde{b}_n - v(\mathbf{1}) + v(\mathbf{0})}{\mathbf{1}^T\hat{A}_n^{-1}\mathbf{1}}\Big). \quad (17)$$

By the same argument used in Section 3.2, this approach estimates a solution to the weighted least squares problem whose optimal solution is the generalized Shapley values $\phi_i(V)$. This adaptation of KernelSHAP is consistent, and the analogous version of unbiased KernelSHAP is consistent and unbiased (see Appendix D). These can be run with our paired sampling approach, and we can also provide uncertainty estimates and detect convergence (Section 4).

6 EXPERIMENTS

We conducted experiments with four datasets to demonstrate the advantages of our Shapley value estimation approach. We used the census income dataset [22], the Portuguese bank marketing dataset [31], the German credit dataset [22], and a breast cancer (BRCA) subtype classification dataset [4]. To avoid overfitting with the BRCA data, we analyzed a random subset of 100 out of 17,814 genes (Appendix E). We trained a LightGBM model [21] for the census data, CatBoost [34] for the credit and bank data, and logistic regression for the BRCA data. Code for our experiments is available online.

To demonstrate local and global explanations with uncertainty estimates, we show examples of SHAP [24] and SAGE [12] values generated using our estimators (Figure 3). Both explanations used a convergence threshold of $t = 0.01$ and display 95% confidence intervals, **which are features not previously offered by KernelSHAP**. We used the dataset sampling approach for both explanations, and for SAGE we used the estimator designed for stochastic cooperative games. These estimators are faster than their unbiased versions, but the results are nearly identical.

Table 1: SHAP estimator run-time comparison. Each value represents the ratio of the average number of samples required relative to the fastest estimator for that dataset (lower is better).

	Census Income	Bank Marketing	German Credit	BRCA Subtypes
Unbiased	380.63	176.45	17437.44	90.40
Unbiased + Paired Sampling	128.60	90.61	422.17	40.44
Original (KernelSHAP)	12.74	7.41	13.74	2.49
Original + Paired Sampling	**1.00**	**1.00**	**1.00**	**1.00**

To measure run-time differences between each estimator when calculating SHAP values, we compared the number of samples required to explain 100 instances for each dataset (Table 1). Rather than reporting the exact number of samples, which is dependent on the convergence threshold, we show the *ratio* between the number of samples required by each estimator; this ratio is independent of the convergence threshold when convergence is defined by the mean squared estimation error falling below a fixed value (Appendix E). Table 1 displays results based on 100 runs for each instance. Results show that the dataset sampling approach (original) is consistently faster than the unbiased estimator, and that paired sampling enables significantly faster convergence. In particular, we find that our paired sampling approach yields a **9× speedup on average over the original KernelSHAP.**

To investigate the accuracy of our convergence forecasting method, we compared the predicted number of samples to the true number across 250 runs. The number of samples depends on the convergence threshold, and we used a threshold $t = 0.005$ for SHAP and $t = 0.02$ for Shapley Effects and SAGE. Figure 4 shows the results for SHAP (using the census data), Shapley Effects (using the bank data) and SAGE (using the BRCA data). In all three cases, **the forecasts become more accurate with more samples, and they vary within an increasingly narrow range around the true number of required samples.** There is a positive bias in the forecast, but the bias diminishes with more samples.

Finally, to demonstrate the speedup from our approach for stochastic cooperative games, we show that our stochastic estimator converges faster than a naive estimator based on the underlying game's expectation (see Section 5.2). We plotted the ratio between the maximum standard deviation and the gap between the smallest and largest values, which we used to detect convergence (using a threshold $t = 0.01$). Figure 5 shows that the stochastic approach dramatically speeds up both SAGE (using the BRCA data) and Shapley Effects (using the bank data), and that the paired sampling technique accelerates convergence for all estimators. The estimators based on the game's expectation are prohibitively slow and could not be run to convergence. **The fastest estimators for both datasets are stochastic estimators using the paired sampling technique,** and only these methods converged for both datasets in the number of samples displayed. As is the case for SHAP, the dataset sampling approach is often faster than the unbiased approach, but the latter is slightly faster for SAGE when using paired sampling.

7 DISCUSSION

This paper described several approaches for estimating Shapley values via linear regression. We first introduced an unbiased version of KernelSHAP, with properties that are simpler to analyze than the original version. We then developed techniques for detecting convergence, calculating uncertainty estimates, and reducing the variance of both the original and unbiased estimators. Finally, we adapted our approach to provide significantly faster estimators for two global explanation methods based on stochastic cooperative games. Our work makes significant strides towards improving the practicality of Shapley value estimation by automatically determining the required number of samples, providing confidence intervals, and accelerating the estimation process.

More broadly, our work contributes to a mature literature on Shapley value estimation [6, 26] and to the growing ML model explanation field [32, 38, 13, 24, 12]. We focused on improving the regression-based approach to Shapley value estimation, and we leave to future work a detailed comparison of this approach to sampling-based [38, 37, 10, 12] and model-specific approximations [1, 25]. We also believe that certain insights from our work may be applicable to LIME, which is based on a similar dataset sampling approach; recent work has noted LIME's high variance when using an insufficient number of samples [3], and an improved understanding of its convergence properties [16, 27] may lead to approaches for automatic convergence detection and uncertainty estimation.

Acknowledgements

This work was funded by the National Science Foundation [CAREER DBI-1552309, and DBI- 1759487]; the American Cancer Society [127332-RSG-15-097-01-TBG]; and the National Institutes of Health [R35 GM 128638, and R01 NIA AG 061132]. We would like to thank Hugh Chen, the Lee Lab, and our AISTATS reviewers for feedback that greatly improved this work.

References

[1] Marco Ancona, Cengiz Öztireli, and Markus Gross. Explaining deep neural networks with a polynomial time algorithm for Shapley values approximation. *arXiv preprint arXiv:1903.10992*, 2019.

[2] Robert JJ Aumann. Economic applications of the Shapley value. In *Game-Theoretic Methods in General Equilibrium Analysis*, pages 121–133. Springer, 1994.

[3] Naman Bansal, Chirag Agarwal, and Anh Nguyen. SAM: The sensitivity of attribution methods to hyperparameters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8673–8683, 2020.

[4] Ashton C Berger, Anil Korkut, Rupa S Kanchi, Apurva M Hegde, Walter Lenoir, Wenbin Liu, Yuexin Liu, Huihui Fan, Hui Shen, Visweswaran Ravikumar, et al. A comprehensive pan-cancer molecular study of gynecologic and breast cancers. *Cancer Cell*, 33(4):690–705, 2018.

[5] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[6] Javier Castro, Daniel Gómez, and Juan Tejada. Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730, 2009.

[7] A Charnes and Daniel Granot. Coalitional and chance-constrained solutions to n-person games. i: The prior satisficing nucleolus. *SIAM Journal on Applied Mathematics*, 31(2):358–367, 1976.

[8] A Charnes, B Golany, M Keane, and J Rousseau. Extremal principle solutions of games in characteristic function form: core, Chebychev and Shapley value generalizations. In *Econometrics of Planning and Efficiency*, pages 123–133. Springer, 1988.

[9] Abraham Charnes and Daniel Granot. Prior solutions: Extensions of convex nucleus solutions to chance-constrained games. Technical report, Texas University at Austin Center for Cybernetic Studies, 1973.

[10] Jianbo Chen, Le Song, Martin J Wainwright, and Michael I Jordan. L-Shapley and C-Shapley: Efficient model interpretation for structured data. *arXiv preprint arXiv:1808.02610*, 2018.

[11] Ian Covert, Scott Lundberg, and Su-In Lee. Explaining by removing: A unified framework for model explanation. *arXiv preprint arXiv:2011.14878*, 2020.

[12] Ian Covert, Scott Lundberg, and Su-In Lee. Understanding global feature contributions with additive importance measures. *Advances in Neural Information Processing Systems*, 34, 2020.

[13] Anupam Datta, Shayak Sen, and Yair Zick. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 598–617. IEEE, 2016.

[14] Guoli Ding, Robert F Lax, Jianhua Chen, and Peter P Chen. Formulas for approximating pseudo-boolean random variables. *Discrete Applied Mathematics*, 156(10):1581–1597, 2008.

[15] Guoli Ding, Robert F Lax, Jianhua Chen, Peter P Chen, and Brian D Marx. Transforms of pseudo-boolean random variables. *Discrete Applied Mathematics*, 158(1):13–24, 2010.

[16] Damien Garreau and Ulrike von Luxburg. Looking deeper into LIME. *arXiv preprint arXiv:2008.11092*, 2020.

[17] Amirata Ghorbani and James Zou. Data Shapley: Equitable valuation of data for machine learning. *arXiv preprint arXiv:1904.02868*, 2019.

[18] Amirata Ghorbani and James Zou. Neuron Shapley: Discovering the responsible neurons. *arXiv preprint arXiv:2002.09815*, 2020.

[19] Michel Grabisch, Jean-Luc Marichal, and Marc Roubens. Equivalent representations of set functions. *Mathematics of Operations Research*, 25(2): 157–178, 2000.

[20] Peter L Hammer and Ron Holzman. Approximations of pseudo-boolean functions; applications to game theory. *Zeitschrift für Operations Research*, 36(1):3–21, 1992.

[21] 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. In *Advances in Neural Information Processing Systems*, pages 3146–3154, 2017.

[22] Moshe Lichman et al. UCI machine learning repository, 2013.

[23] Stan Lipovetsky and Michael Conklin. Analysis of regression in game theory approach. *Applied*

Stochastic Models in Business and Industry, 17 (4):319–330, 2001.

[24] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774, 2017.

[25] 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):2522–5839, 2020.

[26] Sasan Maleki, Long Tran-Thanh, Greg Hines, Talal Rahwan, and Alex Rogers. Bounding the estimation error of sampling-based Shapley value approximation. *arXiv preprint arXiv:1306.4265*, 2013.

[27] Dina Mardaoui and Damien Garreau. An analysis of LIME for text data. *arXiv preprint arXiv:2010.12487*, 2020.

[28] Jean-Luc Marichal and Pierre Mathonet. Weighted Banzhaf power and interaction indexes through weighted approximations of games. *European Journal of Operational Research*, 211 (2):352–358, 2011.

[29] Luke Merrick and Ankur Taly. The explanation game: Explaining machine learning models using Shapley values. *arXiv preprint arXiv:1909.08128*, 2019.

[30] Dov Monderer, Dov Samet, et al. Variations on the Shapley value. *Handbook of Game Theory*, 3: 2055–2076, 2002.

[31] Sérgio Moro, Paulo Cortez, and Paulo Rita. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems*, 62:22–31, 2014.

[32] Art B Owen. Sobol' indices and Shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.

[33] Leon Petrosjan and Georges Zaccour. Time-consistent Shapley value allocation of pollution cost reduction. *Journal of Economic Dynamics and Control*, 27(3):381–398, 2003.

[34] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. Catboost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, pages 6638–6648, 2018.

[35] Marco Tulio Ribeiro, Sameer Singh, and Carlos 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*, pages 1135–1144, 2016.

[36] Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.

[37] Eunhye Song, Barry L Nelson, and Jeremy Staum. Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1060–1083, 2016.

[38] Erik Štrumbelj and Igor Kononenko. Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41(3):647–665, 2014.

[39] Nikola Tarashev, Kostas Tsatsaronis, and Claudio Borio. Risk attribution using the Shapley value: Methodology and policy applications. *Review of Finance*, 20(3):1189–1213, 2016.

[40] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1): 267–288, 1996.

[41] Aad W Van der Vaart. *Asymptotic Statistics*, volume 3. Cambridge University Press, 2000.

[42] BP Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.