
APPROXIMATING THE SHAPLEY VALUE WITHOUT MARGINAL CONTRIBUTIONS

Patrick Kolpaczki

Paderborn University

patrick.kolpaczki@upb.de

Viktor Bengs^{a,b}, Maximilian Muschalik^{a,b}, Eyke Hüllermeier^{a,b}

^aInstitute of Informatics, University of Munich (LMU)

^bMunich Center for Machine Learning

viktor.bengs@lmu.de, maximilian.muschalik@lmu.de, eyke@lmu.de

January 31, 2024

ABSTRACT

The Shapley value, which is arguably the most popular approach for assigning a meaningful contribution value to players in a cooperative game, has recently been used intensively in explainable artificial intelligence. Its meaningfulness is due to axiomatic properties that only the Shapley value satisfies, which, however, comes at the expense of an exact computation growing exponentially with the number of agents. Accordingly, a number of works are devoted to the efficient approximation of the Shapley value, most of them revolve around the notion of an agent’s marginal contribution. In this paper, we propose with *SVARM* and *Stratified SVARM* two parameter-free and domain-independent approximation algorithms based on a representation of the Shapley value detached from the notion of marginal contribution. We prove unmatched theoretical guarantees regarding their approximation quality and provide empirical results including synthetic games as well as common explainability use cases comparing ourselves with state-of-the-art methods.

1 Introduction

Whenever agents can federalize in groups (form coalitions) to accomplish a task and get rewarded with a collective benefit that is to be shared among the group members, the notion of *cooperative game* stemming from game theory is arguably the most favorable concept to model such situations. This is due to its simplicity, which nevertheless allows for covering a whole range of practical applications. The agents are called *players* and are contained in a player set \mathcal{N} . Each possible subset of players $S \subseteq \mathcal{N}$ is understood as a *coalition* and the coalition \mathcal{N} containing all players is called the *grand coalition*. The collective benefit $v(S)$ that a coalition S receives upon formation is given by a *value function* v assigning each coalition a real-valued *worth*.

The connection of cooperative games to (supervised) machine learning is already well-established. The most prominent example is feature importance scores, both local and global, for a machine learning model: features of a dataset can be seen as players, allowing one to interpret a feature subset as a coalition, while the model’s generalization performance using exactly that feature subset is its worth Cohen et al. [2007]. Other applications include evaluating the importance of parameters in a machine learning model, e.g. single neurons in a deep neural network Ghorbani and Zou [2020] or base learners in an ensemble Rozemberczki and Sarkar [2021], or assigning relevance scores to datapoints in a given dataset Ghorbani and Zou [2019]. See Rozemberczki et al. [2022] for a wider overview of its usage in the field of explainable artificial intelligence. Outside the realm of machine learning cooperative games also found applications in operations research Luo et al. [2022], for finding fair compensation mechanisms in electricity grids O’Brien et al. [2015], or even for the purpose of identifying the most influential individuals in terrorist networks van Campen et al. [2018].

In all of these applications, the question naturally arises of how to appropriately determine the contribution of a single player (feature, parameter, etc.) with respect to the grand collective benefit. In other words, how to allocate the worth

$\nu(\mathcal{N})$ of the full player set \mathcal{N} among the players in a fair manner. The indisputably most popular solution to this problem is the *Shapley value* Shapley [1953], which can be intuitively expressed by *marginal contributions*. We call the increase in worth that comes with the inclusion of player i to a coalition S , i.e., the difference $\nu(S \cup \{i\}) - \nu(S)$, player i 's marginal contribution to S . The Shapley value of i is a weighted average of all its marginal contributions to coalitions that do not include i . Its popularity stems from the fact that it is the only solution to satisfy axiomatic properties that arguably capture fairness Shapley [1953].

Despite the appealing theoretical properties of the Shapley value, there is one major drawback with respect to its practical application, as its computational complexity increases exponentially with the number of players n . As a consequence, the exact computation of the Shapley value becomes practically infeasible even for a moderate number of players. This is especially the case where accesses to ν are costly, e.g., re-evaluating a (complex) machine learning model for a specific feature subset, or manipulating training data each time ν is accessed. Recently, several approximation methods have been proposed in search of a remedy, enabling the utilization of the Shapley value in explainable AI (and beyond). However, most works are stiffened towards the notion of marginal contribution, and, consequently, judge algorithms by their achieved approximation accuracy depending on the number of evaluated marginal contributions. This measure does not do justice to the fact that approximations can completely dispense with the consideration of marginal contributions and elicit information from ν in a more efficient way — as we show in this paper. We claim that the number of single accesses to ν should be considered instead, since especially in machine learning, as mentioned above, access to ν is a bottleneck in overall runtime. In this paper, we make up for this deficit by considering the problem of approximating the Shapley values under a fixed *budget* T of evaluations (accesses) of ν .

Contribution. We present a novel representation of the Shapley value that does not rely on the notion of marginal contribution. Our first proposed approximation algorithm *Shapley Value Approximation without Requesting Marginals* (SVARM) exploits this representation and directly samples values of coalitions, facilitating “a swarm of updates”, i.e., multiple Shapley value estimates are updated at once. This is in stark contrast to the usual way of sampling marginal contributions that only allows the update of a single estimate. We prove theoretical guarantees regarding SVARM's precision including the bound of $\mathcal{O}(\frac{\log n}{T-n})$ on its variance.

Based on a partitioning of the set of all coalitions according to their size, we develop with *Stratified SVARM* a refinement of SVARM. The applied stratification materializes a twofold improvement: (i) the homogeneous strata (w.r.t. the coalition worth) significantly accelerate convergence of estimates, (ii) our stratified representation of the Shapley value with decomposed marginal contributions facilitates a mechanism that updates the estimates of *all* players with *each single* coalition sampled. Among other results, we bound its variance by $\mathcal{O}(\frac{\log n}{T-n \log n})$.

Besides our superior theoretical findings, both algorithms possess a number of properties in their favor. More specifically, both are unbiased, parameter-free, incremental, i.e., the available budget has not to be fixed and can be enlarged or cut prematurely, facilitating on-the-fly approximations due to their anytime property, and do not require any knowledge about the latent value function. Moreover, both are domain-independent and not limited to some specific fields, but can be used to approximate the Shapley values of any possible cooperative game.

Finally, we compare our algorithms empirically against other popular competitors, demonstrating their practical usefulness and proving our empirical enhancement *Stratified SVARM*⁺, which samples without replacement to be the first sample-mean-based approach to achieve rivaling state-of-the-art approximation quality. All code including documentation and the technical appendix can be found on GitHub¹.

2 Related Work

The recent rise of explainable AI has incentivized the research on approximation methods for the Shapley value leading to a variety of different algorithms for this purpose. The first distinction to be made is between those that are domain-independent, i.e., able to deal with any cooperative game, and those that are tailored to a specific use case, e.g. assigning Shapley values to single neurons in neural networks, or which impose specific assumptions on the value function. In this paper, we will consider only the former, as it is our goal to provide approximations algorithms independent of the context in which they are applied. The first and so far simplest of this kind is *ApproShapley* Castro et al. [2009], which samples marginal contributions from each player based on randomly drawn permutations of the player set. The variance of each of its Shapley value estimates is bounded by $\mathcal{O}(\frac{n}{T})$. *Stratified Sampling* [Maleki et al., 2013] and *Structured Sampling* [van Campen et al., 2018] both partition the marginal contributions of each player by coalition size in order to stratify the marginal contributions of the population from which to draw a sample, which leads to a variance reduction. While *Stratified Sampling* calculates a sophisticated allocation of samples for each coalition size, *Structured Sampling* simply samples with equal frequencies. Multiple follow-up works suggest

¹<https://github.com/kolpaczki/Approximating-the-Shapley-Value-without-Marginal-Contributions>

specific techniques to improve the sampling allocation over the different coalition sizes [O’Brien et al., 2015, Castro et al., 2017, Burgess and Chapman, 2021].

In order to reduce the variance of the naive sampling approach underlying *ApproShapley*, Illés and Kerényi [2019] suggest to use ergodic sampling, i.e., generating samples that are not independent but still satisfy the strong Law of Large numbers. Quite recently, Mitchell et al. [2022] investigated two techniques for improving *ApproShapley*’s sampling approach. One is based on the theory of reproducing kernel Hilbert spaces, which focuses on minimizing the discrepancies for functions of permutations. The other exploits a geometrical connection between uniform sampling on the Euclidean sphere and uniform sampling over permutations.

Adopting a Bayesian perspective, i.e., by viewing the Shapley values as random variables, Touati et al. [2021] consider approximating the Shapley values by Bayesian estimates (posterior mean, mode, or median), where each posterior distribution of a player’s Shapley value depends on the remaining ones. Utilizing a representation of the Shapley value as an integral Owen [1972], *Owen Sampling* [Okhrati and Lipani, 2020] approximates this integral by sampling marginal contributions using antithetic sampling [Rubinstein and Kroese, 2016, Lomeli et al., 2019] for variance reduction.

A fairly new class of approaches that dissociates itself from the notion of marginal contribution are those that view the Shapley value as a solution of a quadratic program with equality constraints Lundberg and Lee [2017], Simon and Vincent [2020], Covert and Lee [2021]. Another unorthodox approach is to divide the player set into small enough groups for which the Shapley values within these groups can be computed exactly Soufiani et al. [2014], Corder and Decker [2019]. For an overview of approaches related to machine learning we refer to Chen et al. [2023].

3 Problem Statement

The formal notion of a cooperative game is defined by a tuple (\mathcal{N}, ν) consisting of a set of players $\mathcal{N} = \{1, \dots, n\}$ and a value function $\nu : \mathcal{P}(\mathcal{N}) \rightarrow \mathbb{R}$ that assigns to each subset of \mathcal{N} a real-valued number. The value function must satisfy $\nu(\emptyset) = 0$. We call the subsets of \mathcal{N} coalitions, \mathcal{N} itself the grand coalition, and the assigned value $\nu(S)$ to a coalition $S \subseteq \mathcal{N}$ its worth. Given a cooperative game (\mathcal{N}, ν) , the Shapley value assigns each player a share of the grand coalition’s worth. In particular, the Shapley value Shapley [1953] of any player $i \in \mathcal{N}$ is defined as

$$\phi_i = \sum_{S \subseteq \mathcal{N}_i} \frac{1}{n \cdot \binom{n-1}{|S|}} [\nu(S \cup \{i\}) - \nu(S)], \quad (1)$$

where $\mathcal{N}_i := \mathcal{N} \setminus \{i\}$ for each player $i \in \mathcal{N}$. The term $\nu(S \cup \{i\}) - \nu(S)$ is also known as player i ’s marginal contribution to $S \subseteq \mathcal{N}_i$ and captures the increase in collective benefit when player i joins the coalition S . Thus, the Shapley value can be seen as the weighted average of a player’s marginal contributions.

The exact computation of all Shapley values requires the knowledge of the values of all 2^n many coalitions² and is shown to be NP-hard Deng and Papadimitriou [1994]. In light of the exponential computational effort w.r.t. to n , we consider the goal of approximating the Shapley value of all players as precisely as possible for a given *budget* of $T \in \mathbb{N}$ many evaluations (accesses) of ν in discrete time steps $1, \dots, T$. Since $\nu(\emptyset) = 0$ holds by definition, the evaluation of $\nu(\emptyset)$ comes for free without any budget cost. We judge the quality of the estimates $\hat{\phi}_1, \dots, \hat{\phi}_n$ — which are possibly of stochastic nature — obtained by an approximation algorithm after T many evaluations by two criteria that have to be minimized for all $i \in \mathcal{N}$. First, the mean squared error (MSE) of the estimate $\hat{\phi}_i$ is given by

$$\mathbb{E}[(\hat{\phi}_i - \phi_i)^2]. \quad (2)$$

Utilizing the bias-variance decomposition allows us to reduce the squared error to the variance $\mathbb{V}[\hat{\phi}_i]$ of the Shapley value estimate in case that it is unbiased, i.e. $\mathbb{E}[\hat{\phi}_i] = \phi_i$. The second criterion is the probability of $\hat{\phi}_i$ deviating from ϕ_i by more than a fixed $\varepsilon > 0$:

$$\mathbb{P}(|\hat{\phi}_i - \phi_i| > \varepsilon). \quad (3)$$

Both criteria are well-established for measuring the quality of an algorithm approximating the Shapley value.

4 SVARM

Thanks to the distributive law, the formula of the Shapley value for a player i can be rearranged so that it is not its weighted average of marginal contributions, but the difference of the weighted average of coalition values by adding i

²In fact, only $2^n - 1$ many coalitions, as $\nu(\emptyset) = 0$ is known.




Figure 1: Illustration of SVARM's sampling process and update rule: Each player i has two urns $U_i^+ := \{S \cup \{i\} \mid S \subseteq \mathcal{N}_i\}$ and $U_i^- := \{S \mid S \subseteq \mathcal{N}_i\}$ containing marbles which represent coalitions, with mean coalition worth ϕ_i^+ and ϕ_i^- . SVARM alternates between sampling coalitions $A^+ \sim P^+$ and $A^- \sim P^-$. With each drawn coalition all estimates of those urns are updated which contain the corresponding marble. Since each player's two urns form a partition of the powerset $\mathcal{P}(\mathcal{N})$, all players have exactly one urn updated with each sample.

and the weighted average of coalition values without i :

$$\phi_i = \underbrace{\sum_{S \subseteq \mathcal{N}_i} w_S \cdot \nu(S \cup \{i\})}_{=: \phi_i^+} - \underbrace{\sum_{S \subseteq \mathcal{N}_i} w_S \cdot \nu(S)}_{=: \phi_i^-}, \quad (4)$$

with weights $w_S = \frac{1}{n \cdot \binom{n-1}{|S|}}$ for each $S \subseteq \mathcal{N}_i$. We call ϕ_i^+ the positive and ϕ_i^- the negative Shapley value, while we refer to the collective of both as the signed Shapley values. The weighted averages ϕ_i^+ and ϕ_i^- can also be viewed as expected values, i.e., $\phi_i^+ = \mathbb{E}[\nu(S \cup \{i\})]$ and $\phi_i^- = \mathbb{E}[\nu(S)]$, where $S \sim P^w$ and $P^w(S) = w_S$ for all $S \subseteq \mathcal{N}_i$. Note that all weights add up to 1 and thus P^w forms a well-defined probability distribution. In this way, we can approximate each signed Shapley value separately using estimates $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ and combine them into a Shapley value estimate by means of $\hat{\phi}_i = \hat{\phi}_i^+ - \hat{\phi}_i^-$.

In light of this, a naive approach for approximating each signed Shapley value of a player is by sampling some number of M many coalitions $S^{(1)}, \dots, S^{(M)}$ with distribution P^w and using the sample mean as the estimate, i.e., $\hat{\phi}_i^+ = \frac{1}{M} \sum_{m=1}^M \nu(S^{(m)} \cup \{i\})$. However, this would require all $2n$ signed Shapley values (two per player) to be estimated separately by sampling coalitions in a dedicated manner, each of which would lead to an update of only one estimate. This ultimately slows down the convergence of the estimates, especially for large n .

On the basis of the aforementioned representation of the Shapley value, we present the *Shapley Value Approximation without Requesting Marginals* (SVARM) algorithm, a novel approach that updates multiple Shapley value estimates at once with a single evaluation of ν . Its novelty consists of sampling coalitions independently from two specifically chosen distributions P^+ and P^- in an alternating fashion, which allows for a more powerful update rule: each (independently) sampled coalition A^+ from P^+ allows one to update all positive Shapley value estimates $\hat{\phi}_i^+$ of all payers i which are contained in A^+ , i.e., $i \in A^+$. Likewise, for a coalition A^- drawn from P^- , all negative Shapley value estimates $\hat{\phi}_i^-$ for $i \notin A^-$ can be updated.

It is worth noting that, for simplicity, we alternate evenly between the samples from the P^+ and P^- distributions, although one could also use a ratio other than $1/2$. To avoid a bias, both distributions have to be tailored such that the following holds for all $i \in \mathcal{N}$ and $S \subseteq \mathcal{N}_i$:

$$\mathbb{P}(A^+ = S \cup \{i\} \mid i \in A^+) = \mathbb{P}(A^- = S \mid i \notin A^-) = w_S. \quad (5)$$

Algorithm 1 SVARM

Input: $\mathcal{N}, T \in \mathbb{N}$

```

1:  $\hat{\phi}_i^+, \hat{\phi}_i^- \leftarrow 0$  for all  $i \in \mathcal{N}$ 
2:  $c_i^+, c_i^- \leftarrow 1$  for all  $i \in \mathcal{N}$ 
3: WARMUP
4:  $t \leftarrow 2n$ 
5: while  $t + 2 \leq T$  do
6:   Draw  $A^+ \sim P^+$ 
7:   Draw  $A^- \sim P^-$ 
8:    $v^+ \leftarrow \nu(A^+)$ 
9:    $v^- \leftarrow \nu(A^-)$ 
10:  for  $i \in A^+$  do
11:     $\hat{\phi}_i^+ \leftarrow \frac{c_i^+ \hat{\phi}_i^+ + v^+}{c_i^+ + 1}$ 
12:     $c_i^+ \leftarrow c_i^+ + 1$ 
13:  end for
14:  for  $i \in \mathcal{N} \setminus A^-$  do
15:     $\hat{\phi}_i^- \leftarrow \frac{c_i^- \hat{\phi}_i^- + v^-}{c_i^- + 1}$ 
16:     $c_i^- \leftarrow c_i^- + 1$ 
17:  end for
18:   $t \leftarrow t + 2$ 
19: end while
20:  $\hat{\phi}_i \leftarrow \hat{\phi}_i^+ - \hat{\phi}_i^-$  for all  $i \in \mathcal{N}$ 

```

Output: $\hat{\phi}_1, \dots, \hat{\phi}_n$

For this reason, we define the probability distributions over coalitions to sample from as

$$P^+(S) := \frac{1}{|S| \binom{n}{|S|} H_n} \quad \forall S \in \mathcal{P}(\mathcal{N}) \setminus \{\emptyset\}, \quad (6)$$

$$P^-(S) := \frac{1}{(n - |S|) \binom{n}{|S|} H_n} \quad \forall S \in \mathcal{P}(\mathcal{N}) \setminus \{\mathcal{N}\}, \quad (7)$$

where $H_n = \sum_{k=1}^n 1/k$ denotes the n -th harmonic number. Note that both P^+ and P^- assign equal probabilities to coalitions of the same size, so that one can first sample the size and then draw a set uniformly of that size. This pair of distributions is provably the only one to fulfill the required property (see Appendix C.1).

The approach of dividing the Shapley value into two parts and approximating both has already been pursued (although not as formally rigorous) via importance sampling Covert et al. [2019], allowing to update all n estimates with each sample. Wang and Jia [2023] adopt the same representation for the Banzhaf value, and coined the strategy of updating all players' estimates with each sampled coalition the *maximum sample reuse* (MSR) principle. Their approximation algorithm is specifically tailored to the Banzhaf value as it leverages its uniform weights $w_S = \frac{1}{2^{n-1}}$ and is thus, at least not directly, transferable to the Shapley value.

In the following we describe SVARM's procedure with the pseudocode of Algorithm 1. The overall idea of the sampling and update process is illustrated in Figure 1. It starts by initializing the positive and negative Shapley value estimates $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$, and the number of samples c_i^+ and c_i^- collected for each player i . SVARM continues by launching a warm-up phase (see Algorithm 3 in Appendix B). In the main loop, the update rule is applied for as many sampled pairs of coalitions A^+ and A^- as possible until SVARM runs out of budget. In each iteration A^+ is sampled from P^+ and A^- from P^- . The worth of A^+ and A^- is evaluated and stored in v^+ and v^- , requiring two accesses to the value function. The estimate $\hat{\phi}_i^+$ of each player $i \in A^+$ is updated with the worth $\nu(A^+)$ such that $\hat{\phi}_i^+$ is the mean of sampled coalition values. Likewise, the estimate $\hat{\phi}_i^-$ of each player $i \notin A^-$ is updated with the worth $\nu(A^-)$. At the same time, the sample numbers of the respective signed Shapley value estimates are also updated. Finally, SVARM computes its Shapley value estimate $\hat{\phi}_i$ of ϕ_i for each i according to Equation (4). Note that since only the quantities $\hat{\phi}_i^+, \hat{\phi}_i^-, c_i^+$, and c_i^- are stored for each player, its space complexity is in $\mathcal{O}(n)$. Moreover, SVARM is incremental and can be stopped at any time to return its estimates after executing line 20, or it can be run further with increased budget.

Theoretical analysis. In the following we present theoretical results for SVARM. All proofs are given in Appendix C of the technical appendix. For the remainder of this section we assume that a minimum budget of $T \geq 2n + 2$ is given. This assumption guarantees the completion of the warm-up phase such that each positive and negative Shapley value estimate has at least one sample and an additional pair sampled in the loop. The lower bound on T is essentially twice the number of players n , which is a fairly weak assumption. We denote by $\bar{T} := T - 2n$ the number of time steps (budget) left after the warm-up phase. Moreover, we assume \bar{T} to be even for sake of simplicity such that a lower bound on the number of sampled pairs in the main part can be expressed by $\frac{\bar{T}}{2} - n$. We begin with the unbiasedness of the estimates maintained by SVARM allowing us later to reduce the mean squared error (MSE) of each estimate to its variance.

Theorem 1. *The Shapley value estimate $\hat{\phi}_i$ of any $i \in \mathcal{N}$ obtained by SVARM is unbiased, i.e.,*

$$\mathbb{E}[\hat{\phi}_i] = \phi_i.$$

Next, we give a bound on the variance of each Shapley value estimate. For this purpose, we introduce notation for the variances of coalition values contained in ϕ_i^+ and ϕ_i^- . For a random set $A_i \subseteq \mathcal{N}_i$ distributed according to P^w let

$$\sigma_i^{+2} := \mathbb{V}[\nu(A_i \cup \{i\})] \text{ and } \sigma_i^{-2} := \mathbb{V}[\nu(A_i)] . \quad (8)$$

Theorem 2. *The variance of any player's Shapley value estimate $\hat{\phi}_i$ obtained by SVARM is bounded by*

$$\mathbb{V}[\hat{\phi}_i] \leq \frac{2H_n}{\bar{T}}(\sigma_i^{+2} + \sigma_i^{-2}).$$

Combining the unbiasedness in Theorem 1 with the latter variance bound implies the following result on the MSE.

Corollary 1. *The MSE of any player's Shapley value estimate $\hat{\phi}_i$ obtained by SVARM is bounded by*

$$\mathbb{E}[(\hat{\phi}_i - \phi_i)^2] \leq \frac{2H_n}{\bar{T}}(\sigma_i^{+2} + \sigma_i^{-2}).$$

Assuming that each variance term σ_i^{+2} and σ_i^{-2} is bounded by some constant independent of n (and T), the MSE bound in Corollary 1 is in $\mathcal{O}(\frac{\log n}{\bar{T}-n})$ and so is the variance bound in Theorem 2. Note that this assumption is rather mild and satisfied if the underlying value function is bounded by constants independent of n , which again is the case for a wide range of games and in particular in explainable AI for global and local feature importance based on classification probabilities lying between 0 and 1. Further, as T is growing linearly with n by assumption, the denominator is essentially driven by the asymptotics of T . Thus, the dependency on n is logarithmic, which is a significant improvement over existing theoretical results having a linear dependency on n like $\mathcal{O}(\frac{n}{T})$ for *ApproShapley* [Castro et al., 2009] or possibly worse [Simon and Vincent, 2020]. Finally, we present two probabilistic bounds on the approximated Shapley value. The first utilizes the variance bound shown in Theorem 2 by applying Chebyshev's inequality.

Theorem 3. *The probability that the Shapley value estimate $\hat{\phi}_i$ of any fixed player $i \in \mathcal{N}$ deviates from ϕ_i by a margin of any fixed $\varepsilon > 0$ or greater is bounded by*

$$\mathbb{P}(|\hat{\phi}_i - \phi_i| \geq \varepsilon) \leq \frac{2H_n}{\varepsilon^2 \bar{T}}(\sigma_i^{-2} + \sigma_i^{+2}).$$

The presented bound is in $\mathcal{O}(\frac{\log n}{\bar{T}-n})$ and improves upon the bound derived by Chebyshev's inequality of $\mathcal{O}(\frac{n}{T})$ for *ApproShapley* Maleki et al. [2013]. Our second bound derived by Hoeffding's inequality is tighter, but requires the introduction of notation for the ranges of $\nu(A_i)$ and $\nu(A_i \cup \{i\})$:

$$r_i^+ := \max_{S \subseteq \mathcal{N}_i} \nu(S \cup \{i\}) - \min_{S \subseteq \mathcal{N}_i} \nu(S \cup \{i\}), \quad (9)$$

$$r_i^- := \max_{S \subseteq \mathcal{N}_i} \nu(S) - \min_{S \subseteq \mathcal{N}_i} \nu(S). \quad (10)$$

Theorem 4. *The probability that the Shapley value estimate $\hat{\phi}_i$ of any fixed player $i \in \mathcal{N}$ deviates from ϕ_i by a margin of any fixed $\varepsilon > 0$ or greater is bounded by*

$$\mathbb{P}(|\hat{\phi}_i - \phi_i| \geq \varepsilon) \leq 2e^{-\frac{\varepsilon^2}{4H_n}} + 4\frac{e^{-\Psi \lfloor \frac{\bar{T}}{4H_n} \rfloor}}{e^\Psi - 1},$$

where $\Psi = 2\varepsilon^2/(r_i^+ + r_i^-)^2$.

Note that this bound is exponentially decreasing with T and can be expressed asymptotically as $\mathcal{O}(e^{-\frac{T-n}{(\log n)^2}})$. In comparison, the bounds of $\mathcal{O}(e^{-\frac{T}{n}})$ for *ApproShapley*, $\mathcal{O}(ne^{-\frac{T}{n^3}})$ for *Stratified Sampling* Maleki et al. [2013], and the projected SGD variant Simon and Vincent [2020] show worse asymptotic dependencies on n in comparison.

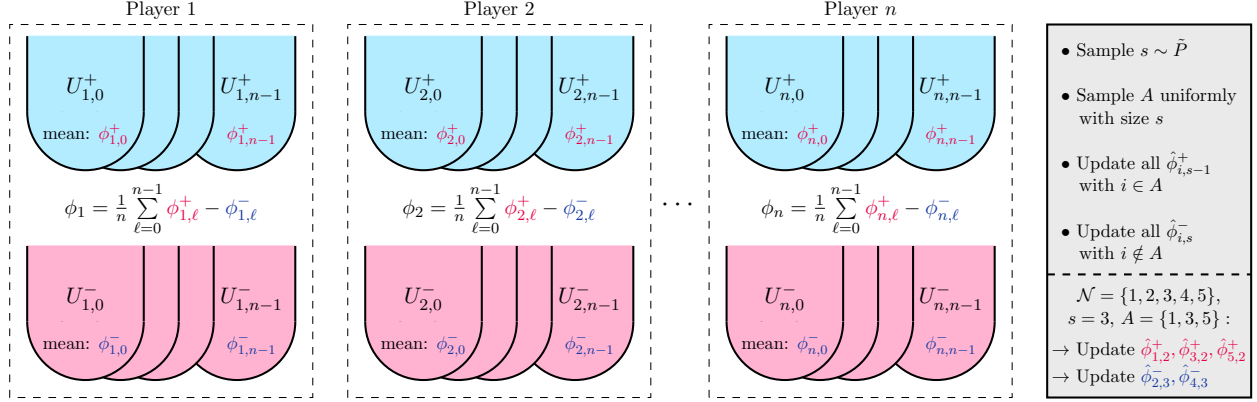


Figure 2: Illustration of Stratified SVARM's sampling process and update rule: Each player i has urns $U_{i,\ell}^+ := \{S \cup \{i\} \mid S \subseteq \mathcal{N}_i, |S| = \ell\}$ and $U_{i,\ell}^- := \{S \mid S \subseteq \mathcal{N}_i, |S| = \ell\}$ for all $\ell \in \{0, \dots, n-1\}$, $2n$ in total, containing marbles which represent coalitions, with mean coalition worth $\phi_{i,\ell}^+$ and $\phi_{i,\ell}^-$. Stratified SVARM samples in each time step t a coalition $A_t \subseteq \mathcal{N}$ and updates the estimates of all players' urns that contain the corresponding marble. Since each player's urns form a partition of the powerset $\mathcal{P}(\mathcal{N})$, all players have exactly one urn updated with each sample.

5 Stratified SVARM

On the basis of the representation of the Shapley value in Equation (4), we develop another approximation algorithm named *Stratified SVARM* to further pursue and reach the maximum sample reuse principle. Its crux is a refinement of *SVARM* obtained by stratifying the positive and the negative Shapley value ϕ_i^+ and ϕ_i^- . We exploit the latter to develop an even more powerful update rule that allows for updating all players simultaneously with each single coalition sampled. Both, ϕ_i^+ and ϕ_i^- can be rewritten using stratification such that each becomes an average of strata, whereas the strata themselves are averages of the coalitions' worth:

$$\phi_i^+ = \frac{1}{n} \sum_{\ell=0}^{n-1} \frac{1}{\binom{n-1}{\ell}} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S|=\ell}} \nu(S \cup \{i\}) =: \frac{1}{n} \sum_{\ell=0}^{n-1} \phi_{i,\ell}^+, \quad (11)$$

$$\phi_i^- = \frac{1}{n} \sum_{\ell=0}^{n-1} \frac{1}{\binom{n-1}{\ell}} \sum_{\substack{S \subseteq \mathcal{N}_i \\ |S|=\ell}} \nu(S) =: \frac{1}{n} \sum_{\ell=0}^{n-1} \phi_{i,\ell}^-. \quad (12)$$

We call $\phi_{i,\ell}^+$ the ℓ -th positive Shapley subvalue and $\phi_{i,\ell}^-$ the ℓ -th negative Shapley subvalue for all $\ell \in \mathcal{L} := \{0, \dots, n-1\}$. Now, we can write ϕ_i as

$$\phi_i = \frac{1}{n} \sum_{\ell=0}^{n-1} \phi_{i,\ell}^+ - \phi_{i,\ell}^-. \quad (13)$$

Note that this representation of ϕ_i coincides with Equation 6 in Ancona et al. [2019]. Intuitively speaking at the example of ϕ_i^+ (and analogously for ϕ_i^-), we partition the population of coalitions contained in ϕ_i^+ into n strata. Each *stratum* $\phi_{i,\ell}^+$ comprises all coalitions which include the player i and have cardinality $\ell+1$. Instead of sampling directly for ϕ_i^+ , the stratification allows one to sample coalitions from each stratum, obtain mean estimates $\hat{\phi}_{i,\ell}^+$, and aggregate them to

$$\hat{\phi}_i^+ = \frac{1}{n} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^+ \quad (14)$$

in order to obtain an estimate for ϕ_i^+ . Due to the increase in homogeneity of the strata in comparison to their origin population, caused by the shared size and inclusion or exclusion of i for coalitions in the same stratum, one would expect the strata to have significantly lower variances and ranges resulting in approximations of better quality compared to *SVARM*. In combination with our bounds shown in Theorem 2 and Theorem 4, this should result in approximations of better quality. In the following we present further techniques for improvement which we apply for Stratified *SVARM* (Algorithm 2).

Algorithm 2 Stratified SVARM

Input: $\mathcal{N}, T \in \mathbb{N}$

- 1: $\hat{\phi}_{i,\ell}^+, \hat{\phi}_{i,\ell}^- \leftarrow 0$ for all $i \in \mathcal{N}$ and $\ell \in \mathcal{L}$
- 2: $c_{i,\ell}^+, c_{i,\ell}^- \leftarrow 0$ for all $i \in \mathcal{N}$ and $\ell \in \mathcal{L}$
- 3: **EXACTCALCULATION**(\mathcal{N})
- 4: **WARMUP**⁺(\mathcal{N})
- 5: **WARMUP**⁻(\mathcal{N})
- 6: $t \leftarrow 2n + 1 + 2 \sum_{s=2}^{n-2} \lceil \frac{n}{s} \rceil$
- 7: **while** $t < T$ **do**
- 8: Draw $s_t \sim \tilde{P}$
- 9: Draw A_t from $\{S \subseteq \mathcal{N} \mid |S| = s_t\}$ uniformly
- 10: **UPDATE**(A_t)
- 11: $t \leftarrow t + 1$
- 12: **end while**
- 13: $\hat{\phi}_i \leftarrow \frac{1}{n} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^+ - \hat{\phi}_{i,\ell}^-$ for all $i \in \mathcal{N}$

Output: $\hat{\phi}_1, \dots, \hat{\phi}_n$

Exact calculation. First, we observe that some strata contain very few coalitions. Thus, we calculate $\phi_{i,0}^+, \phi_{i,n-2}^+, \phi_{i,n-1}^+, \phi_{i,1}^-$, and $\phi_{i,n-1}^-$ for all players exactly by evaluating ν for all coalitions of size 1, $n-1$, and n . This requires $2n+1$ many evaluations of ν (see Algorithm 5 in Appendix B). We already obtain $\phi_{i,0}^- = \nu(\emptyset) = 0$ by definition. As a consequence, we can exclude the sizes 0, 1, $n-1$, and n from further consideration. We assume for the remainder that $n \geq 4$, otherwise we would have already calculated all Shapley values exactly.

Refined warm-up. Next, we split the warm-up into two parts, one for the positive, the other for the negative Shapley subvalues (see Algorithm 6 and 7 in Appendix B). Each collects for each estimate $\hat{\phi}_{i,\ell}^+$ or $\hat{\phi}_{i,\ell}^-$, respectively, one sample and consumes a budget of $\sum_{s=2}^{n-2} \lceil \frac{n}{s} \rceil$.

Enhanced update rule. Thanks to the stratified representation of the Shapley value, we can enhance SVARM's update rule and update with each sampled coalition $A_t \subseteq \mathcal{N}$ the estimates $\hat{\phi}_{i,|A_t|-1}^+$ for all $i \in A_t$ and $\hat{\phi}_{i,|A_t|}^-$ for all $i \notin A_t$. Thus, we can update all estimates $\hat{\phi}_i$ at once with a single sample. This enhanced update step is given in Algorithm 4 (see Appendix B) and illustrated in Figure 2. In order to obtain unbiased estimates, it suffices to select an arbitrary size s of the coalition A to be sampled and draw A uniformly at random from the set of coalitions with size s . We go one step further and choose not only the coalition A , but also the size s randomly according to a specifically tailored probability distribution \tilde{P} over $\{2, \dots, n-2\}$, which leads to simpler bounds in our theoretical analysis in which each stratum receives the same weight. We define for n even:

$$\tilde{P}(s) := \begin{cases} \frac{n \log n - 1}{2sn \log n \left(H_{\frac{n}{2}-1} - 1 \right)} & \text{if } s \leq \frac{n-2}{2} \\ \frac{1}{n \log n} & \text{if } s = \frac{n}{2} \\ \frac{n \log n - 1}{2(n-s)n \log n \left(H_{\frac{n}{2}-1} - 1 \right)} & \text{otherwise} \end{cases} \quad \text{and for } n \text{ odd: } \tilde{P}(s) := \begin{cases} \frac{1}{2s \left(H_{\frac{n-1}{2}-1} \right)} & \text{if } s \leq \frac{n-1}{2} \\ \frac{1}{2(n-s) \left(H_{\frac{n-1}{2}-1} \right)} & \text{otherwise} \end{cases} \quad (15)$$

Note that Stratified SVARM is incremental just as SVARM, but in contrast, requires quadratic space $\mathcal{O}(n^2)$ as it stores estimates and counters for each player *and* stratum.

Theoretical analysis. Similar to SVARM, we present in the following our theoretical results for Stratified SVARM. All proofs are given in Appendix D. Again, we assume a minimum budget of $T \geq 2n + 1 + 2 \sum_{s=2}^{n-2} \lceil \frac{n}{s} \rceil =: W \in \mathcal{O}(n \log n)$, guaranteeing the completion of the warm-up phase, and denote by $\bar{T} = T - W$ the budget left after the warm-up phase. We start by showing that Stratified SVARM is not afflicted with any bias.

Theorem 5. *The Shapley value estimate $\hat{\phi}_i$ of any $i \in \mathcal{N}$ obtained by Stratified SVARM is unbiased, i.e.,*

$$\mathbb{E}[\hat{\phi}_i] = \phi_i.$$

Next, we consider the variance of the Shapley value estimates and quickly introduce some notation. Let $A_{i,\ell} \subseteq \mathcal{N}_i$ be a random coalition of size ℓ distributed with $\mathbb{P}(A_{i,\ell} = S) = \binom{n-1}{\ell}^{-1}$. Define the strata variances

$$\sigma_{i,\ell}^{+2} := \mathbb{V}[\nu(A_{i,\ell} \cup \{i\})] \text{ and } \sigma_{i,\ell}^{-2} := \mathbb{V}[\nu(A_{i,\ell})]. \quad (16)$$

Theorem 6. *The variance of any player's Shapley value estimate $\hat{\phi}_i$ obtained by Stratified SVARM is bounded by*

$$\mathbb{V}[\hat{\phi}_i] \leq \frac{2 \log n}{nT} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

Together with the unbiasedness shown in Theorem 5, the variance bound implies the following MSE bound.

Corollary 2. *The MSE of any player's Shapley value estimate $\hat{\phi}_i$ obtained by Stratified SVARM is bounded by*

$$\mathbb{E}[(\hat{\phi}_i - \phi_i)^2] \leq \frac{2 \log n}{nT} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

With our choice of the sampling distribution \tilde{P} we achieved an easily interpretable bound on the MSE in which each stratum variance is equally weighted. Assuming that each stratum variance is bounded by some constant independent of n , the MSE bound in Corollary 2 is in $\mathcal{O}(\frac{\log n}{T - n \log n})$. Note that, by assumption, T is growing log-linearly with n so that the denominator is essentially driven by the asymptotics of T . Again, compared to existing theoretical results, with linear dependence on n , the logarithmic dependence on n is a significant improvement. Still, it is worth emphasizing that the more homogeneous strata with lower variances constitute the core improvement of Stratified SVARM, which are not reflected within the O-notation. Our first probabilistic bound is obtained by Chebyshev's inequality and the bound from Theorem 6.

Theorem 7. *The probability that the Shapley value estimate $\hat{\phi}_i$ of any fixed player $i \in \mathcal{N}$ deviates from ϕ_i by a margin of any fixed $\varepsilon > 0$ or greater is bounded by*

$$\mathbb{P}(|\hat{\phi}_i - \phi_i| \geq \varepsilon) \leq \frac{2 \log n}{\varepsilon^2 n T} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

Lastly, our second probabilistic bound derived via Hoeffding's inequality is tighter, but less trivial. It requires some further notation, namely the ranges of the strata values:

$$r_{i,\ell}^+ := \max_{S \subseteq \mathcal{N}_i: |S|=\ell} \nu(S \cup \{i\}) - \min_{S \subseteq \mathcal{N}_i: |S|=\ell} \nu(S \cup \{i\}), \quad (17)$$

$$r_{i,\ell}^- := \max_{S \subseteq \mathcal{N}_i: |S|=\ell} \nu(S) - \min_{S \subseteq \mathcal{N}_i: |S|=\ell} \nu(S). \quad (18)$$

Theorem 8. *The probability that the Shapley value estimate $\hat{\phi}_i$ of any fixed player $i \in \mathcal{N}$ deviates from ϕ_i by a margin of any fixed $\varepsilon > 0$ or greater is bounded by*

$$\mathbb{P}(|\hat{\phi}_i - \phi_i| \geq \varepsilon) \leq 2(n-3) \left(e^{-\frac{\bar{T}}{8n^2(\log n)^2}} + 2 \frac{e^{-\Psi \lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{e^\Psi - 1} \right),$$

where $\Psi = 2\varepsilon^2 n^2 / \left(\sum_{\ell=1}^{n-3} r_{i,\ell}^+ + r_{i,\ell+1}^- \right)^2$.

This bound is of order $\mathcal{O}(ne^{-\frac{T-n \log n}{n^2(\log n)^2}})$ showing a slightly worse dependency on n compared to Theorem 4 due to the introduction of strata.

6 Empirical Results

To complement our theoretical findings, we evaluate our algorithms and its competitors on commonly considered synthetic cooperative games and explainable AI scenarios in which Shapley values need to be approximated. In particular, we select parameterless algorithms that do not rely on provided knowledge about the value function of the problem instance at hand, since ours do not either. Besides the sampling distribution \tilde{P} over coalition sizes proposed for Stratified SVARM (S-SVARM), we also consider sampling with the simpler uniform distribution over all sizes from

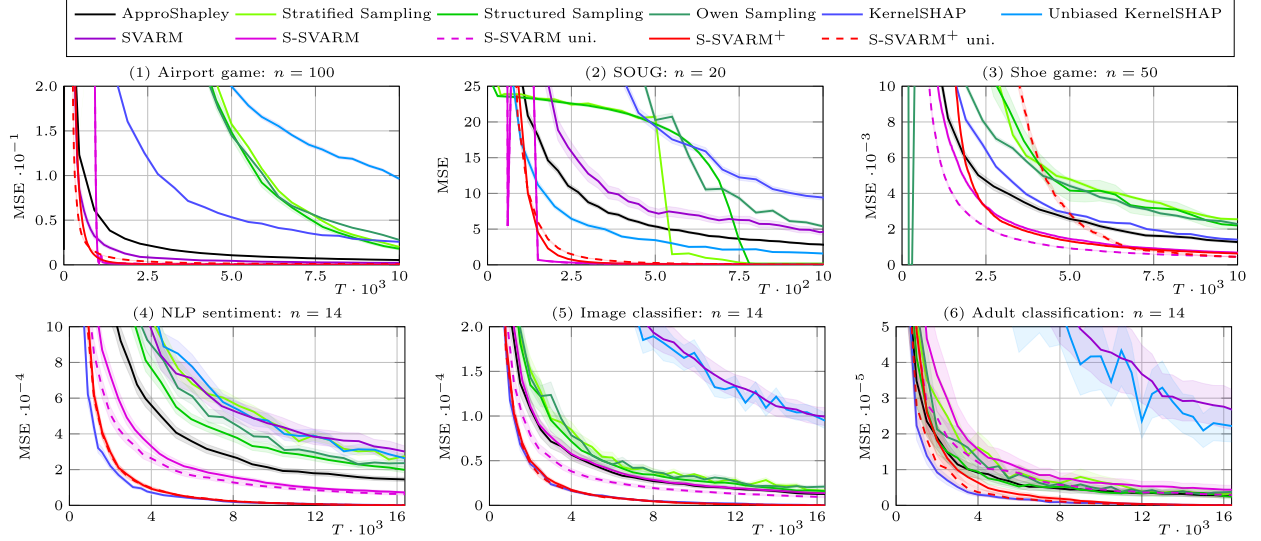


Figure 3: Averaged MSE and standard errors over 100 repetitions in dependence of fixed budget T : (1) Airport game, (2) Shoe game, (3) SOUG game, (4) NLP sentiment analysis, (5) Image classifier, (6) Adult classification.

2 to $n - 2$ (S-SVARM uniform). In order to allow for a fair comparison with KernelSHAP, which samples coalitions without replacement, we include with S-SVARM⁺ (uniform) an empirical version of S-SVARM without the warm-up that also samples without replacement to compensate for this underlying advantage (see Algorithm 8 in Appendix B), which obviously comes at the price of space complexity linear in T .

We run the algorithms multiple times on the selected game types and measure their performances by the mean squared error (MSE) averaged over all players and runs depending on a range of fixed budget values T . Measuring the approximation quality by the MSE requires the true Shapley values of the considered games to be available. These are either given by a polynomial closed-form solution for the synthetic games (see Section 6.1) or we compute them exhaustively for our explanation tasks (see Section 6.2). The results of our evaluation are shown in Figure 3 and are presented in more detail in Appendix F.

As already said, we judge the algorithms’ approximation qualities in dependence on the spent budget (model evaluations) T instead of the consumed runtime. In fact, the algorithms differ in actual runtime. For example SVARM performs less arithmetic operations than Stratified SVARM since it does not update all players’ estimates $\hat{\phi}_i^+$ or $\hat{\phi}_i^-$ with each sample. Some algorithms, e.g. KernelSHAP, vary strongly in their time consumption per sample since a costly quadratic optimization problem needs to be solved after observing all samples. We intentionally avoid the runtime comparison for three reasons: (i) the observed runtimes may differ depending on the actual implementation, (ii) the fixed-budget setting facilitates a coherent theoretical analysis where the observed information is restricted, (iii) evaluating the worth of a coalition poses the bottleneck in explanation tasks, rendering the difference in performed arithmetic operations negligible.

6.1 Synthetic games

Cooperative games with polynomial closed-form solutions of their Shapley values are well suited for tracking the approximation error for large player numbers. We exploit this fact and investigate a broad range of player numbers n which are significantly higher than those for the explanation tasks. We conduct experiments on the predefined Shoe and Airport game as done in Castro et al. [2009, 2017]. Their degree of non-additivity poses a difficult challenge to all approximation algorithms. Further, we consider randomly generated Sum of Unanimity Games (SOUG) games van Campen et al. [2018] which are capable of representing any cooperative game. The value function and Shapley values of each game are given in Appendix E.

We observe that S-SVARM itself already shows reliably good approximation performance across all considered games and budget ranges. It is significantly superior to its competitors AppoShapley and KernelSHAP and as expected, S-SVARM⁺ extends the lead in approximation quality even more. In contrast, SVARM can rarely keep up with its refined counterpart S-SVARM. However, in light of the bounds on the MSEs in Corollary 1 and 2 this is not surprising: SVARM’s MSE bound scales linearly with the variances σ_i^{+2} and σ_i^{-2} of *all* coalition values containing respectively

not containing i , while the relevant variance terms $\sigma_{i,\ell}^{+2}$ and $\sigma_{i,\ell}^{-2}$ for S-SVARM are restricted to coalitions of fixed size. In most games, the latter terms are significantly lower since coalitions of the same size are plausibly closer in worth. Finally, S-SVARM is quite robust regarding the magnitude of the standard errors.

6.2 Explainability games

We further conduct experiments on cooperative games stemming from real-world explainability scenarios, in particular, use cases in which local feature importance of machine learning models are to be quantified via Shapley values. The NLP sentiment analysis game is based on the DistilBERT Sanh et al. [2019] model architecture and consists of randomly selected movie reviews from the IMDB dataset Maas et al. [2011] containing 14 words. Missing features are masked in the tokenized representation and the value of a set is its sentiment score. In the image classifier game, we explain the output of a ResNet18 He et al. [2016] trained on ImageNet Deng et al. [2009]. The images' pixels are summarized into $n = 14$ super-pixels and absent features are masked with mean imputation. The worth of a coalition is the returned class probability of the model (using only the present super-pixels) for the class of the original prediction which was made with all pixels being present. For the adult classification game, we train a gradient-boosted tree model on the adult dataset Becker and Kohavi [1996]. A coalition's worth is the predicted class probability of the true income class (income above or below 50 000) of the given datapoint with the absent features being removed via mean imputation. Since no polynomial closed-form solution exists for the Shapley values in these games, we compute them exhaustively, limiting us to a feasible number of players for which we can track the MSE. While this restricts us to a player number (tokens, superpixels, features) of $n = 14$ due to limited computational resources, this is arguably still an appropriate and commonly appearing number of entities involved in an explanation task. We refer to Appendix E for a more detailed explanation of the chosen games.

A first observation is the close head-to-head race between S-SVARM⁺ and KernelSHAP across the considered games leaving all other methods behind. Thus, S-SVARM⁺ is the first sample-mean-based approach achieving rivaling state-of-the-art approximation quality. KernelSHAP's counterpart Unbiased KernelSHAP, designed to facilitate approximation guarantees similar to our theoretical results which KernelSHAP lacks, is clearly outperformed by S-SVARM. Given the consistency demonstrated by S-SVARM and S-SVARM⁺, we claim that both constitute a reliable choice under absence of domain knowledge. We conjecture that the reason for the slight performance decrease of S-SVARM from synthetic to explainability games lies not only within the latent structure of ν , but is also caused by the lower player numbers. As our theoretical results indicate, its sample efficiency grows with n due to its enhanced update rule. However, conducting experiments with larger n becomes computationally prohibitive for explainability games, since the Shapley values have to be calculated exhaustively in order to track the approximation error. Further, our results indicate the robustness of S-SVARM⁽⁺⁾ w.r.t. the utilized distribution \tilde{P} , which allows us to use the uniform distribution without performance loss, and secondly shows that our derived distribution is not just a theoretical artifact, but a valid contribution to express simpler bounds which are easier to grasp and interpret.

7 Conclusion

We considered the problem of precisely approximating the Shapley value of all players in a cooperative game under the restriction that the value function can be evaluated only a given number of times. We presented a reformulation of the Shapley value, detached from the ubiquitous notion of marginal contribution, facilitating the approximation by estimates of which a multitude can be updated with each access to the value function. On this basis, we proposed two approximation algorithms, SVARM and Stratified SVARM, which have a number of desirable properties. Both are parameter-free, incremental, domain-independent, unbiased, and do not require any prior knowledge of the value function. Further, Stratified SVARM shows a satisfying compromise between peak approximation quality and consistency across all considered games, paired with unmatched theoretical guarantees regarding its approximation quality. While fulfilling more desirable properties and not having to solve a quadratic optimization problem of size T in comparison to the state-of-the-art method KernelSHAP, effectively disabling on-the-fly approximations, our simpler sample-mean-based method Stratified SVARM⁺ can fully keep up in common explainable AI scenarios, and even shows empirical superiority on synthetic games.

Limitations and future work. The quadratically growing number of strata w.r.t. n might pose a challenge for higher player numbers, which future work could remedy by applying a coarser stratification that assigns multiple coalition sizes to a single stratum. One could investigate the empirical behavior in further popular explanation domains such as data valuation, federated learning, or neuron importance and extend our evaluation to scenarios with higher player numbers. Since the true Shapley values are not accessible for larger n , a different measure of approximation quality than the MSE needs to be taken for reference. The convergence speed of the estimates is a naturally arising alternative. Our empirical results give further evidence for the non-existence of a universally best approximation

algorithm and encourage future research into the cause of the observed differences in performance w.r.t. the game type. Further, it would be interesting to analyze whether structural properties of the value function, such as monotonicity or submodularity, have an impact on the approximation quality of both algorithms.

Acknowledgments

This research was supported by the research training group Dataninja (Trustworthy AI for Seamless Problem Solving: Next Generation Intelligence Joins Robust Data Analysis) funded by the German federal state of North Rhine-Westphalia. We gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): TRR 318/1 2021 – 438445824. We would like to thank Fabian Fumagalli and especially Patrick Becker for their efforts in supporting our implementation.

References

- Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurélien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012. doi: 10.1109/TPAMI.2012.120.
- Marco Ancona, Cengiz Öztireli, and Markus H. Gross. Explaining deep neural networks with a polynomial time algorithm for shapley value approximation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 272–281, 2019.
- Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996.
- Mark Alexander Burgess and Archie C. Chapman. Approximating the shapley value using stratified empirical bernstein sampling. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 73–81, 2021.
- Javier Castro, Daniel Gómez, and Juan Tejada. Polynomial calculation of the shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730, 2009.
- Javier Castro, Daniel Gómez, Elisenda Molina, and Juan Tejada. Improving polynomial estimation of the shapley value by stratified random sampling with optimum allocation. *Computers & Operations Research*, 82:180–188, 2017.
- M. T. Chao and W. E. Strawderman. Negative moments of positive random variables. *Journal of the American Statistical Association*, 67(338):429–431, 1972.
- Hugh Chen, Ian C. Covert, Scott M. Lundberg, and Su-In Lee. Algorithms to estimate shapley value feature attributions. *Nature Machine Intelligence*, 5:590–601, 2023.
- Shay B. Cohen, Gideon Dror, and Eytan Ruppin. Feature selection via coalitional game theory. *Neural Computation*, 19(7):1939–1961, 2007.
- Kevin Corder and Keith Decker. Shapley value approximation with divisive clustering. In *18th IEEE International Conference On Machine Learning And Applications*, pages 234–239, 2019.
- Ian Covert and Su-In Lee. Improving kernelshap: Practical shapley value estimation using linear regression. In *24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *PMLR*, pages 3457–3465, 2021.
- Ian Covert, Scott Lundberg, and Su-In Lee. Shapley feature utility. In *Machine Learning in Computational Biology*, 2019.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.
- Amirata Ghorbani and James Y. Zou. Data shapley: Equitable valuation of data for machine learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2242–2251, 2019.
- Amirata Ghorbani and James Y. Zou. Neuron shapley: Discovering the responsible neurons. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

- Ferenc Illés and Péter Kerényi. Estimation of the shapley value by ergodic sampling. *CoRR*, abs/1906.05224, 2019.
- Maria Lomeli, Mark Rowland, Arthur Gretton, and Zoubin Ghahramani. Antithetic and monte carlo kernel estimators for partial rankings. *Statistics and Computing*, 29(5):1127–1147, 2019.
- Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4768–4777, 2017.
- Chunlin Luo, Xiaoyang Zhou, and Benjamin Lev. Core, shapley value, nucleolus and nash bargaining solution: A survey of recent developments and applications in operations management. *Omega*, 110:102638, 2022. ISSN 0305-0483.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, 2011.
- Sasan Maleki, Long Tran-Thanh, Greg Hines, Talal Rahwan, and Alex Rogers. Bounding the estimation error of sampling-based shapley value approximation with/without stratifying. *CoRR*, abs/1306.4265, 2013.
- Rory Mitchell, Joshua Cooper, Eibe Frank, and Geoffrey Holmes. Sampling permutations for shapley value estimation. *Journal of Machine Learning Research*, 23(43):1–46, 2022.
- Gearóid O’Brien, Abbas El Gamal, and Ram Rajagopal. Shapley value estimation for compensation of participants in demand response programs. *IEEE Transactions on Smart Grid*, 6(6):2837–2844, 2015.
- Ramin Okhrati and Aldo Lipani. A multilinear sampling algorithm to estimate shapley values. In *25th International Conference on Pattern Recognition*, pages 7992–7999, 2020.
- Guillermo Owen. Multilinear extensions of games. *Management Science*, 18:64–79, 1972.
- Benedek Rozemberczki and Rik Sarkar. The shapley value of classifiers in ensemble games. In *30th ACM International Conference on Information and Knowledge Management*, pages 1558–1567, 2021.
- Benedek Rozemberczki, Lauren Watson, Péter Bayer, Hao-Tsung Yang, Oliver Kiss, Sebastian Nilsson, and Rik Sarkar. The shapley value in machine learning. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, pages 5572–5579, 2022.
- Reuven Y Rubinstein and Dirk P Kroese. *Simulation and the Monte Carlo Method*. John Wiley & Sons, 2016.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108, 2019.
- L. S. Shapley. A value for n-person games. In *Contributions to the Theory of Games, Volume II*, pages 307–318. Princeton University Press, 1953.
- Grah Simon and Thouvenot Vincent. A projected stochastic gradient algorithm for estimating shapley value applied in attribute importance. In *Machine Learning and Knowledge Extraction*, pages 97–115, 2020.
- Hossein Azari Soufiani, David Maxwell Chickering, Denis Xavier Charles, and David C. Parkes. Approximating the Shapley Value via Multi-Issue Decompositions. In *Proceedings of the International conference on Autonomous Agents and Multi-Agent Systems*, volume 2, 2014.
- Sofiane Touati, Mohammed Said Radjef, and Lakhdar Sais. A bayesian monte carlo method for computing the shapley value: Application to weighted voting and bin packing games. *Computers & Operations Research*, 125:105094, 2021. doi: 10.1016/j.cor.2020.105094.
- Tjeerd van Campen, Herbert Hamers, Bart Husslage, and Roy Lindelauf. A new approximation method for the shapley value applied to the wtc 9/11 terrorist attack. *Social Network Analysis and Mining*, 8(3):1–12, 2018.
- Jiachen T. Wang and Ruoxi Jia. Data banzhaf: A robust data valuation framework for machine learning. In *26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *PMLR*, pages 6388–6421, 2023.

A List of Symbols

Table 1: List of frequently symbols used throughout the paper.

Problem setting	
\mathcal{N}	set of players
\mathcal{N}_i	set of players without i
n	number of players
ν	value function
T	budget, number of allowed evaluations of ν
ϕ_i	Shapley value of player i
$\hat{\phi}_i$	estimated Shapley value of player i
SVARM	
ϕ_i^+	positive Shapley value
ϕ_i^-	negative Shapley value
$\hat{\phi}_i^+$	estimated positive Shapley value
$\hat{\phi}_i^-$	estimated negative Shapley value
P^+	sampling probability distribution over coalitions to estimate ϕ_i^+
P^-	sampling probability distribution over coalitions to estimate ϕ_i^-
\bar{T}	remaining budget after completion of the warm-up phase
σ_i^{+2}	variance of coalition values including player i
σ_i^{-2}	variance of coalition values excluding player i
r_i^+	range of coalition values including player i
r_i^-	range of coalition values excluding player i
\bar{m}_i^+	number of sampled coalitions after the warm-up phase to update $\hat{\phi}_i^+$
\bar{m}_i^-	number of sampled coalitions after the warm-up phase to update $\hat{\phi}_i^-$
m_i^+	total number of sampled coalitions to update $\hat{\phi}_i^+$
m_i^-	total number of sampled coalitions to update $\hat{\phi}_i^-$
Stratified SVARM	
$\phi_{i,\ell}^+$	ℓ -th positive Shapley subvalue
$\phi_{i,\ell}^-$	ℓ -th negative Shapley subvalue
$\hat{\phi}_{i,\ell}^+$	estimated ℓ -th positive Shapley subvalue
$\hat{\phi}_{i,\ell}^-$	estimated ℓ -th negative Shapley subvalue
\bar{P}	sampling probability distribution over coalition sizes
\bar{T}	remaining budget after completion of the warm-up phase
$\sigma_{i,\ell}^{+2}$	variance of coalition values in the ℓ -th stratum including player i
$\sigma_{i,\ell}^{-2}$	variance of coalition values in the ℓ -th stratum excluding player i
$r_{i,\ell}^+$	range of coalition values in the ℓ -th stratum including player i
$r_{i,\ell}^-$	range of coalition values in the ℓ -th stratum excluding player i
$\bar{m}_{i,\ell}^+$	number of sampled coalitions after the warm-up phase to update $\hat{\phi}_{i,\ell}^+$
$\bar{m}_{i,\ell}^-$	number of sampled coalitions after the warm-up phase to update $\hat{\phi}_{i,\ell}^-$
$m_{i,\ell}^+$	total number of sampled coalitions to update $\hat{\phi}_{i,\ell}^+$
$m_{i,\ell}^-$	total number of sampled coalitions to update $\hat{\phi}_{i,\ell}^-$

B Further Pseudocode

B.1 SVARM

Algorithm 3 WARMUP

```

1: for  $i \in \mathcal{N}$  do
2:   Draw  $A^+$  and  $A^-$  i.i.d. from  $P^w$ 
3:    $\hat{\phi}_i^+ \leftarrow \nu(A^+ \cup \{i\})$ 
4:    $\hat{\phi}_i^- \leftarrow \nu(A^-)$ 
5: end for
    
```

The warm-up of SVARM samples for each player i two coalitions A^+ and A^- , both drawn i.i.d. according to the weights w_S , i.e., the probability distribution P^w , and updates $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$, which needs a budget of $2n$ in total. This ensures that each estimate is based on at least one sample.

B.2 Stratified SVARM

Algorithm 4 UPDATE(A)

```

1:  $v \leftarrow \nu(A)$ 
2: for  $i \in A$  do
3:    $\hat{\phi}_{i,|A|-1}^+ \leftarrow \frac{c_{i,|A|-1}^+ \cdot \hat{\phi}_{i,|A|-1}^+ + v}{c_{i,|A|-1}^+ + 1}$ 
4:    $c_{i,|A|-1}^+ \leftarrow c_{i,|A|-1}^+ + 1$ 
5: end for
6: for  $i \in \mathcal{N} \setminus A$  do
7:    $\hat{\phi}_{i,|A|}^- \leftarrow \frac{c_{i,|A|}^- \cdot \hat{\phi}_{i,|A|}^- + v}{c_{i,|A|}^- + 1}$ 
8:    $c_{i,|A|}^- \leftarrow c_{i,|A|}^- + 1$ 
9: end for
    
```

Stratified SVARM's update procedure updates exactly one Shapley subvalue of each player given a coalition A . It consumes only one budget token by storing the worth of A in the variable v (line 1). The first loop increments for all players $i \in A$ their counter $c_{i,|A|-1}^+$ by 1 and updates the $|A| - 1$ -th positive Shapley subvalue estimate $\hat{\phi}_{i,|A|-1}^+$ to be the average over all values of coalitions which are contained in that stratum of player i . Analogously, the second loop increments for all players i not contained in A their counter $c_{i,|A|}^-$ by 1 and updates the $|A|$ -th negative Shapley subvalue estimate $\hat{\phi}_{i,|A|}^-$ to be the average over all values of coalitions which are contained in that stratum of player i .

Algorithm 5 EXACTCALCULATION(\mathcal{N})

```

1: for  $s \in \{1, n-1, n\}$  do
2:   for  $A \in \{S \subseteq \mathcal{N} \mid |S| = s\}$  do
3:     UPDATE( $A$ )
4:   end for
5: end for
    
```

The exact calculation evaluates all coalitions of size $1, n-1$ and n , thus $2n+1$ in total. For each coalition, the update procedure is called. Effectively, this leads to exactly computed strata $\hat{\phi}_{i,0}^+ = \phi_{i,0}^+, \hat{\phi}_{i,n-2}^+ = \phi_{i,n-2}^+, \hat{\phi}_{i,n-1}^+ = \phi_{i,n-1}^+, \hat{\phi}_{i,1}^- = \phi_{i,1}^-, \hat{\phi}_{i,n-1}^- = \phi_{i,n-1}^-$ and counters $c_{i,0}^+ = 1, c_{i,n-2}^+ = n-1, c_{i,n-1}^+ = 1, c_{i,1}^- = n-1, c_{i,n-1}^- = 1$.

Algorithm 6 $\text{WARMUP}^+(\mathcal{N})$

```

1: for  $s = 2, \dots, n - 2$  do
2:   Draw a permutation  $\pi$  of  $\mathcal{N}$  u.a.r.
3:   for  $k = 0, \dots, \lfloor \frac{n}{s} \rfloor - 1$  do
4:      $A \leftarrow \{\pi(1 + ks), \dots, \pi(s + ks)\}$ 
5:      $v \leftarrow \nu(A)$ 
6:     for  $i \in A$  do
7:        $\hat{\phi}_{i,s-1}^+ \leftarrow v$ 
8:        $c_{i,s-1}^+ \leftarrow 1$ 
9:     end for
10:  end for
11:  if  $n \bmod s \neq 0$  then
12:     $A \leftarrow \{\pi(n - (n \bmod s) + 1), \dots, \pi(n)\}$ 
13:    Draw  $B \in \{S \subseteq \mathcal{N} \setminus A \mid |S| = s - (n \bmod s)\}$ 
14:     $v \leftarrow \nu(A \cup B)$ 
15:    for  $i \in A$  do
16:       $\hat{\phi}_{i,s-1}^+ \leftarrow v$ 
17:       $c_{i,s-1}^+ \leftarrow 1$ 
18:    end for
19:  end if
20: end for

```

The warm-up for the positive Shapley subvalues iterates over all coalition sizes from 2 to $n - 2$ (line 1) and draws for each size s a permutation π of \mathcal{N} uniformly at random (line 2). The ordering π is sliced into coalitions of size s and each of them is used to update only the players contained in that particular coalition (lines 6–9). In particular, since each coalition A is the first to be observed for the corresponding players' stratum, the estimate $\hat{\phi}_{i,s-1}^+$ is set to the worth of A and its counter $c_{i,s-1}^+$ is set to 1. Note that for each coalition A of size s only one access to ν is made to update all affected s many players. In case that n is not a multiple of s , some players less than s are left over at the end of π (line 11). We group those with other random players to form a coalition of size s , but only update with the worth of that coalition the left out players (lines 15–18). Note that the warm-up comes without any bias, since for each player i and stratum estimate $\hat{\phi}_{i,s}$ each coalition $A \subseteq \mathcal{N}$ with $i \in A$ and $|A| = s$ has the same probability of being chosen. Finally, $\lceil \frac{n}{s} \rceil$ many coalitions are evaluated for each size s , resulting in $\sum_{s=2}^{n-2} \lceil \frac{n}{s} \rceil \in \mathcal{O}(n \log n)$ total evaluations.

Algorithm 7 $\text{WARMUP}^-(\mathcal{N})$

```

1: for  $s = 2, \dots, n - 2$  do
2:   Draw a permutation  $\pi$  of  $\mathcal{N}$  u.a.r.
3:   for  $k = 0, \dots, \lfloor \frac{n}{s} \rfloor - 1$  do
4:      $A \leftarrow \{\pi(1 + ks), \dots, \pi(s + ks)\}$ 
5:      $v \leftarrow \nu(\mathcal{N} \setminus A)$ 
6:     for  $i \in A$  do
7:        $\hat{\phi}_{i,n-s}^- \leftarrow v$ 
8:        $c_{i,n-s}^- \leftarrow 1$ 
9:     end for
10:  end for
11:  if  $n \bmod s \neq 0$  then
12:     $A \leftarrow \{\pi(n - (n \bmod s) + 1), \dots, \pi(n)\}$ 
13:    Draw  $B \in \{S \subseteq \mathcal{N} \setminus A \mid |S| = s - (n \bmod s)\}$  u.a.r.
14:     $v \leftarrow \nu(\mathcal{N} \setminus (A \cup B))$ 
15:    for  $i \in A$  do
16:       $\hat{\phi}_{i,n-s}^- \leftarrow v$ 
17:       $c_{i,n-s}^- \leftarrow 1$ 
18:    end for
19:  end if
20: end for

```

The warm-up for the negative subvalues proceeds analogously to the previously presented positive warm-up. Instead of $\hat{\phi}_{i,s-1}^+$ and $c_{i,s-1}^+$, $\hat{\phi}_{i,n-s}^-$ and $c_{i,n-s}^-$ are updated with the wort of $\mathcal{N} \setminus A$ for all players contained in the coalition A .

B.3 Stratified SVARM⁺

Algorithm 8 Stratified SVARM⁺

Input: $\mathcal{N}, T \in \mathbb{N}$

- 1: $\hat{\phi}_{i,\ell}^+, \hat{\phi}_{i,\ell}^- \leftarrow 0$ for all $i \in \mathcal{N}$ and $\ell \in \mathcal{L}$
- 2: $c_{i,\ell}^+, c_{i,\ell}^- \leftarrow 0$ for all $i \in \mathcal{N}$ and $\ell \in \mathcal{L}$
- 3: **EXACTCALCULATION**(\mathcal{N})
- 4: $t \leftarrow 2n + 1$ {consumed budget}
- 5: $w_s \leftarrow \tilde{P}(s)$ for all $s \in \{2, \dots, n-2\}$ {sampling weight of size s }
- 6: $L_s \leftarrow \emptyset$ for all $s \in \{2, \dots, n-2\}$ {sampled coalitions of size s }
- 7: $m_s \leftarrow \binom{n}{s}$ for all $s \in \{2, \dots, n-2\}$ {reamining sets to sample of size s }
- 8: **while** $t < T$ and $m_s > 0$ for some s **do**
- 9: Draw $s_t \in \{2, \dots, n-2\}$ with probability $\frac{w_{s_t} m_{s_t}}{\sum_{s'=2}^{n-2} w_{s'} m_{s'}}$
- 10: Draw A_t from $\{S \subseteq \mathcal{N} \mid |S| = s_t, S \notin L_{s_t}\}$ u.a.r.
- 11: **UPDATE**(A_t)
- 12: $m_{s_t} \leftarrow m_{s_t} - 1$
- 13: $t \leftarrow t + 1$
- 14: **end while**
- 15: $\hat{\phi}_i \leftarrow \frac{1}{|\{c_{i,\ell}^+ | \ell \in \{0, \dots, n-1\}\}|} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^+ - \frac{1}{|\{c_{i,\ell}^- | \ell \in \{0, \dots, n-1\}\}|} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^-$ for all $i \in \mathcal{N}$

Output: $\hat{\phi}_1, \dots, \hat{\phi}_n$

Stratified SVARM⁺ is a modification of Stratified SVARM to deliver better empirical performance with only two slight changes that do not alter the method on a conceptual level. First, we remove the warm-up since it is less efficient in the sense that not all players estimates are updated with each sampled coalition. Hence, we only consume a budget of $2n + 1$ due to the exact calculation of the border strata before entering the main loop. Although it is extremely unlikely for a sufficiently large chosen budget T and an appropriate distribution \tilde{P} over the coalition sizes, it can happen that some $c_{i,\ell}^+$ or $c_{i,\ell}^-$ are zero. In this case dividing by n the total number of strata per sign per player in line 15 would cause an unnecessary bias. Instead, we average only over all strata for which at least one sample has been observed, i.e. $c_{i,\ell}^+ > 0$ respectively $c_{i,\ell}^- > 0$. Second and most important, we sample coalitions without replacement. We are aware that different ways of implementing this exist (saving substantial amounts of runtime), but we choose to demonstrate it as simply as possible. Effectively each coalition of size $s \in \{2, \dots, n-2\}$ is assigned the weight $\frac{\tilde{P}(s)}{\binom{n}{s}}$, such that coalitions of the same size have the same weight and their weight sums up to $\tilde{P}(s)$. In each time step t a remaining coalition A_t is drawn with probability proportional to its weight (its own weight divided by the sum of all remaining coalitions' weights). We realize this by a two-step procedure: first the size s_t is drawn in line 9, then a remaining coalition of size s_t is drawn uniformly at random in line 10. For this purpose we keep track of all so far sampled coalitions of a given size s in L_s (line 6) and the number of coalitions m_s of size s left to sample (line 7). Finally, we added the condition that at least one coalition must be left to sample to the loop in line 8, in case that T is chosen larger than $2^n - 1$.

C SVARM Analysis

Notation:

- Let $\bar{T} = T - 2n$.
- Let A_i be a random set with $\mathbb{P}(A_i = S) = w_S$ for all $S \subseteq \mathcal{N}_i$.
- Let $\sigma_i^{+2} = \mathbb{V}[\nu(A_i \cup \{i\})]$.
- Let $\sigma_i^{-2} = \mathbb{V}[\nu(A_i)]$.
- Let \bar{m}_i^+ be number of sampled coalitions A^+ after the warm-up phase that contain i .
- Let \bar{m}_i^- be number of sampled coalitions A^- after the warm-up phase that do not contain i .
- Let $m_i^+ = \bar{m}_i^+ + 1$ be total number of samples for $\hat{\phi}_i^+$.
- Let $m_i^- = \bar{m}_i^- + 1$ be total number of samples for $\hat{\phi}_i^-$.
- Let $r_i^+ = \max_{S \subseteq \mathcal{N}_i} \nu(S \cup \{i\}) - \min_{S \subseteq \mathcal{N}_i} \nu(S \cup \{i\})$ be the range of $\nu(A_i \cup \{i\})$.
- Let $r_i^- = \max_{S \subseteq \mathcal{N}_i} \nu(S) - \min_{S \subseteq \mathcal{N}_i} \nu(S)$ be the range of $\nu(A_i)$.

Assumptions:

- \bar{T} is even
- $\bar{T} > 0$

C.1 Unbiasedness of Shapley Value Estimates

To start with, we prove that the distributions P^+ and P^- are well-defined.

Lemma 1. *The distributions P^+ and P^- over $\mathcal{P}(\mathcal{N})$ are well-defined, i.e.,*

$$\sum_{S \subseteq \mathcal{N}} P^+(S) = \sum_{S \subseteq \mathcal{N}} P^-(S) = 1.$$

Proof. The statement is easily shown for P^+ by grouping the coalitions by size. We derive:

$$\begin{aligned} & \sum_{S \subseteq \mathcal{N}} P^+(S) \\ &= \sum_{\ell=1}^n \sum_{\substack{S \subseteq \mathcal{N} \\ |S|=\ell}} P^+(S) \\ &= \sum_{\ell=1}^n \sum_{\substack{S \subseteq \mathcal{N} \\ |S|=\ell}} \frac{1}{\ell \binom{n}{\ell} H_n} \\ &= \frac{1}{H_n} \sum_{\ell=1}^n \frac{1}{\ell} \\ &= 1. \end{aligned}$$

One can prove the desired property analogously for P^- . □

For the remainder of this section we assume that $T \geq 2n + 1$ such that the warm-up phase can be completed by SVARM.

Lemma 2. *For each player $i \in \mathcal{N}$ the positive and negative Shapley Value estimates $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ are unbiased, i.e.,*

$$\mathbb{E}[\hat{\phi}_i^+] = \phi_i^+ \quad \text{and} \quad \mathbb{E}[\hat{\phi}_i^-] = \phi_i^-.$$

Proof. Let \bar{m}_i^+ be the number of coalitions sampled after the warm-up phase that contain i and m_i^+ be the total number of samples used to update $\hat{\phi}_i^+$, thus $m_i^+ = \bar{m}_i^+ + 1$. Further, let A_m^+ for $m \in \{1, 3, 5, \dots, T-1\}$ be the sampled coalitions for updating the positive Shapley values $(\hat{\phi}_i^+)_{i \in \mathcal{N}}$, then we can write the positive Shapley value of player $i \in \mathcal{N}$ as

$$\begin{aligned}\hat{\phi}_i^+ &= \frac{1}{m_i^+} \sum_{\tilde{m}=1}^{T/2} \nu(A_{2\tilde{m}-1}^+) \mathbb{I}_{\{i \in A_{2\tilde{m}-1}^+\}} \\ &= \frac{1}{m_i^+} \left(\nu(A_{2i-1}^+) + \sum_{\tilde{m}=n}^{T/2} \nu(A_{2\tilde{m}-1}^+) \mathbb{I}_{\{i \in A_{2\tilde{m}-1}^+\}} \right),\end{aligned}\tag{19}$$

where \mathbb{I} denotes the indicator function, and we used that during the warm-up phase ($m \leq 2n$) there is for each player i only one A_m^+ to update the corresponding positive Shapley value, namely at time step $2i-1$. First, we show for each odd $m \geq 2n$ and $S \subseteq \mathcal{N}_i$ that $\mathbb{P}(A_m^+ = S \cup \{i\} \mid i \in A_m^+) = w_S$. Note that since $A_m^+ \sim P^+$ (see (6)) it holds that

$$\begin{aligned}\mathbb{P}(i \in A_m^+) &= \sum_{\ell=1}^n \mathbb{P}(i \in A_m^+, |A_m^+| = \ell) \\ &= \sum_{\ell=1}^n \mathbb{P}(i \in A_m^+ \mid |A_m^+| = \ell) \cdot \mathbb{P}(|A_m^+| = \ell) \\ &= \sum_{\ell=1}^n \frac{\ell}{n} \cdot \frac{1}{\ell \cdot H_{n-1}} \\ &= \frac{1}{H_n}.\end{aligned}\tag{20}$$

With this, we derive

$$\begin{aligned}\mathbb{P}(A_m^+ = S \cup \{i\} \mid i \in A_m^+) &= \frac{\mathbb{P}(A_m^+ = S \cup \{i\})}{\mathbb{P}(i \in A_m^+)} \\ &= H_n \cdot P^+(S \cup \{i\}) \\ &= \frac{1}{(|S+1|) \cdot \binom{n}{|S+1|}} \\ &= w_S.\end{aligned}$$

Since $A_{2i-1}^+ \setminus \{i\} \sim P^w$ it holds that $\mathbb{P}(A_{2i-1}^+ \setminus \{i\} = S) = w_S$ for any $i \in \mathcal{N}$ and $S \subset \mathcal{N}_i$.

Taking all of this into account, we derive for $\hat{\phi}_i^+$ using (19):

$$\begin{aligned}\mathbb{E}[\hat{\phi}_i^+ \mid m_i^+] &= \frac{1}{m_i^+} \left(\mathbb{E}[\nu(A_{2i-1}^+) \mid m_i^+] + \mathbb{E} \left[\sum_{\tilde{m}=n}^{T/2} \nu(A_{2\tilde{m}-1}^+) \mathbb{I}_{\{i \in A_{2\tilde{m}-1}^+\}} \mid m_i^+ \right] \right) \\ &= \frac{1}{m_i^+} \left(\sum_{S \subseteq \mathcal{N}_i} \mathbb{P}(A_{2i-1}^+ \setminus \{i\} = S) \cdot \nu(S \cup \{i\}) \right. \\ &\quad \left. + \mathbb{E} \left[\sum_{\tilde{m}=n}^{T/2} \mathbb{I}_{\{i \in A_{2\tilde{m}-1}^+\}} \sum_{S \subset \mathcal{N}_i} \mathbb{P}(A_{2\tilde{m}-1}^+ = S \cup \{i\} \mid i \in A_{2\tilde{m}-1}^+) \cdot \nu(S \cup \{i\}) \mid m_i^+ \right] \right) \\ &= \frac{1}{m_i^+} \left(\sum_{S \subseteq \mathcal{N}_i} w_S \cdot \nu(S \cup \{i\}) + \bar{m}_i^+ \sum_{S \subseteq \mathcal{N}_i} w_S \cdot \nu(S \cup \{i\}) \right) \\ &= \phi_i^+.\end{aligned}$$

Finally, we conclude:

$$\begin{aligned}\mathbb{E}[\hat{\phi}_i^+] &= \sum_{m=1}^{\frac{\bar{T}}{2}} \mathbb{E}[\hat{\phi}_i^+ | m_i^+ = m] \cdot \mathbb{P}(m_i^+ = m) \\ &= \sum_{m=1}^{\frac{\bar{T}}{2}} \phi_i^+ \cdot \mathbb{P}(m_i^+ = m) \\ &= \phi_i^+.\end{aligned}$$

Analogously we derive $\mathbb{E}[\hat{\phi}_i^-] = \phi_i^-$ by defining \bar{m}_i^- , m_i^- , and A_m^- for $m \in \{2, 4, 6, \dots, T\}$ similarly as for their positive counterparts. \square

Theorem 1 For each player $i \in \mathcal{N}$ the estimate $\hat{\phi}_i$ obtained by SVARM is unbiased, i.e.,

$$\mathbb{E}[\hat{\phi}_i] = \phi_i.$$

Proof. We apply Lemma 2 and obtain in combination with Equation (4):

$$\begin{aligned}\mathbb{E}[\hat{\phi}_i] &= \mathbb{E}[\hat{\phi}_i^+] - \mathbb{E}[\hat{\phi}_i^-] \\ &= \phi_i^+ - \phi_i^- \\ &= \phi_i.\end{aligned}$$

\square

C.2 Sample Numbers

Lemma 3. For any $i \in \mathcal{N}$ the expected number of updates of $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ after the warm-up phase is

$$\mathbb{E}[\bar{m}_i^+] = \mathbb{E}[\bar{m}_i^-] = \frac{\bar{T}}{2H_n}.$$

Proof. First, we observe that \bar{m}_i^+ is binomially distributed with $\bar{m}_i^+ \sim \text{Bin}\left(\frac{\bar{T}}{2}, \frac{1}{H_n}\right)$ because $\frac{\bar{T}}{2}$ many pairs are sampled and each independently sampled coalition A^+ contains the player i with probability H_n^{-1} , see (20).

Consequently, we obtain

$$\mathbb{E}[\bar{m}_i^+] = \frac{\bar{T}}{2H_n}.$$

Similarly, we observe that \bar{m}_i^- is also binomially distributed with $\bar{m}_i^- \sim \text{Bin}\left(\frac{\bar{T}}{2}, \frac{1}{H_n}\right)$, leading to the same expected number of updates. \square

C.3 Variance and Squared Error

Lemma 4. The variance of any player's Shapley value estimate $\hat{\phi}_i$ given the number of samples m_i^+ and m_i^- is exactly

$$\mathbb{V}[\hat{\phi}_i | m_i^+, m_i^-] = \frac{\sigma_i^{+2}}{m_i^+} + \frac{\sigma_i^{-2}}{m_i^-}.$$

Proof. We first decompose the variance of $\hat{\phi}_i$ into the variances of $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ and their covariance:

$$\mathbb{V}[\hat{\phi}_i | m_i^+, m_i^-] = \left(\mathbb{V}[\hat{\phi}_i^+ | m_i^+] + \mathbb{V}[\hat{\phi}_i^- | m_i^-] - 2\text{Cov}(\hat{\phi}_i^+, \hat{\phi}_i^- | m_i^+, m_i^-) \right)$$

We derive for the variance of $\hat{\phi}_i^+$:

$$\begin{aligned}\mathbb{V}[\hat{\phi}_i^+ | m_i^+] &= \frac{1}{m_i^{+2}} \sum_{m=0}^{\bar{m}_i^+} \mathbb{V}[\nu(A_{i,m}^+)] \\ &= \frac{\sigma_i^{+2}}{m_i^+}.\end{aligned}$$

Similarly we obtain for $\hat{\phi}_i^-$:

$$\mathbb{V}[\hat{\phi}_i^- | m_i^-] = \frac{\sigma_i^{-2}}{m_i^-}.$$

The covariance of $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ is zero because both are updated with sampled coalitions drawn independently of each other. Thus, we conclude:

$$\mathbb{V}[\hat{\phi}_i | m_i^+, m_i^-] = \frac{\sigma_i^{+2}}{m_i^+} + \frac{\sigma_i^{-2}}{m_i^-}.$$

□

Lemma 5. *For the sample numbers of any player $i \in \mathcal{N}$ holds*

$$\mathbb{E}\left[\frac{1}{m_i^+}\right] = \mathbb{E}\left[\frac{1}{m_i^-}\right] \leq \frac{2H_n}{T}.$$

Proof. By combining Equation (3.4) in Chao and Strawderman [1972]:

$$\mathbb{E}\left[\frac{1}{1+X}\right] = \frac{1 - (1-p)^{m+1}}{(m+1)p} \leq \frac{1}{mp} = \frac{1}{\mathbb{E}[X]},$$

for any binomially distributed random variable $X \sim \text{Bin}(m, p)$ with Lemma 3, we obtain:

$$\mathbb{E}\left[\frac{1}{m_i^+}\right] = \mathbb{E}\left[\frac{1}{1 + \bar{m}_i^+}\right] \leq \frac{1}{\mathbb{E}[\bar{m}_i^+]} = \frac{2H_n}{T}.$$

Notice that m_i^+ and m_i^- are identically distributed.

□

Theorem 2 *The variance of any player's Shapley value estimate $\hat{\phi}_i$ is bounded by*

$$\mathbb{V}[\hat{\phi}_i] \leq \frac{2H_n}{T} (\sigma_i^{+2} + \sigma_i^{-2}).$$

Proof. The combination of Lemma 4 and Lemma 5 yields:

$$\begin{aligned}\mathbb{V}[\hat{\phi}_i] &= \mathbb{E}\left[\mathbb{V}[\hat{\phi}_i | m_i^+, m_i^-]\right] \\ &= \mathbb{E}\left[\frac{\sigma_i^{+2}}{m_i^+} + \frac{\sigma_i^{-2}}{m_i^-}\right] \\ &= \sigma_i^{+2} \mathbb{E}\left[\frac{1}{m_i^+}\right] + \sigma_i^{-2} \mathbb{E}\left[\frac{1}{m_i^-}\right] \\ &\leq \frac{2H_n}{T} (\sigma_i^{+2} + \sigma_i^{-2}).\end{aligned}$$

□

Corollary 1 *The expected squared error of any player's Shapley value estimate is bounded by*

$$\mathbb{E}\left[\left(\hat{\phi}_i - \phi_i\right)^2\right] \leq \frac{2H_n}{T} (\sigma_i^{+2} + \sigma_i^{-2}).$$

Proof. The bias-variance decomposition allows us to plug in the unbiasedness of $\hat{\phi}_i$ shown in Theorem 1 and the bound on the variance from

$$\begin{aligned}\mathbb{E} \left[\left(\hat{\phi}_i - \phi_i \right)^2 \right] &= \left(\mathbb{E} \left[\hat{\phi}_i \right] - \phi_i \right)^2 + \mathbb{V} \left[\hat{\phi}_i \right] \\ &\leq \frac{2H_n}{T} \left(\sigma_i^{+2} + \sigma_i^{-2} \right).\end{aligned}$$

□

C.4 Probabilistic Bounds

Theorem 3 Fix any player $i \in \mathcal{N}$ and $\varepsilon > 0$. The probability that the Shapley value estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater is bounded by

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \frac{2H_n}{\varepsilon^2 T} \left(\sigma_i^{+2} + \sigma_i^{-2} \right).$$

Proof. The bound on the variance of $\hat{\phi}_i$ in Theorem 2 allows us to apply Chebyshev's inequality:

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \frac{\mathbb{V} \left[\hat{\phi}_i \right]}{\varepsilon^2} \leq \frac{2H_n}{\varepsilon^2 T} \left(\sigma_i^{+2} + \sigma_i^{-2} \right).$$

□

Corollary 3. Fix any player $i \in \mathcal{N}$ and $\delta \in (0, 1]$. The Shapley value estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater with probability not greater than δ , i.e.,

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \delta \quad \text{for} \quad \varepsilon = \sqrt{\frac{2H_n}{\delta T} \left(\sigma_i^{+2} + \sigma_i^{-2} \right)}.$$

Lemma 6. For any fixed player $i \in \mathcal{N}$ and $\varepsilon > 0$ holds

$$\mathbb{P} \left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ \right) \leq 2 \exp \left(-\frac{2m_i^+ \varepsilon^2}{r_i^{+2}} \right) \quad \text{and} \quad \mathbb{P} \left(|\hat{\phi}_i^- - \phi_i^-| \geq \varepsilon \mid m_i^- \right) \leq 2 \exp \left(-\frac{2m_i^- \varepsilon^2}{r_i^{-2}} \right).$$

Proof. We prove the statement for $\hat{\phi}_i^+$ by making use of Hoeffding's inequality in combination with the unbiasedness of the positive and negative Shapley value estimates shown in Lemma 2. The proof for $\hat{\phi}_i^-$ is analogous.

$$\begin{aligned}\mathbb{P} \left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ \right) &= \mathbb{P} \left(\left| \hat{\phi}_i^+ - \mathbb{E} \left[\hat{\phi}_i^+ \right] \right| \geq m_i^+ \varepsilon \mid m_i^+ \right) \\ &= \mathbb{P} \left(\left| \sum_{m=0}^{m_i^+} \nu(A_{i,m}^+) - \mathbb{E} \left[\sum_{m=0}^{m_i^+} \nu(A_{i,m}^+) \right] \right| \geq m_i^+ \varepsilon \mid m_i^+ \right) \\ &\leq 2 \exp \left(-\frac{2m_i^+ \varepsilon^2}{r_i^{+2}} \right).\end{aligned}$$

□

Lemma 7. For any fixed player $i \in \mathcal{N}$ and $\varepsilon > 0$ holds:

$$\begin{aligned}\bullet \quad \mathbb{P} \left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \right) &\leq \exp \left(-\frac{\bar{T}}{4H_n^2} \right) + 2 \frac{\exp \left(-\frac{2\varepsilon^2}{r_i^{+2}} \right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp \left(\frac{2\varepsilon^2}{r_i^{+2}} \right) - 1}, \\ \bullet \quad \mathbb{P} \left(|\hat{\phi}_i^- - \phi_i^-| \geq \varepsilon \right) &\leq \exp \left(-\frac{\bar{T}}{4H_n^2} \right) + 2 \frac{\exp \left(-\frac{2\varepsilon^2}{r_i^{-2}} \right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp \left(\frac{2\varepsilon^2}{r_i^{-2}} \right) - 1}.\end{aligned}$$

Proof. We prove the statement for $\hat{\phi}_i^+$. The proof for $\hat{\phi}_i^-$ is analogous. To begin with, we derive with the help of Hoeffding's inequality for binomial distributions a bound for the probability of \bar{m}_i^+ not exceeding $\frac{\bar{T}}{4H_n}$:

$$\begin{aligned} \mathbb{P}\left(\bar{m}_i^+ \leq \frac{\bar{T}}{4H_n}\right) &\leq \mathbb{P}\left(\mathbb{E}[\bar{m}_i^+] - \bar{m}_i^+ \geq \mathbb{E}[\bar{m}_i^+] - \frac{\bar{T}}{4H_n}\right) \\ &\leq \exp\left(-\frac{4\left(\mathbb{E}[\bar{m}_i^+] - \frac{\bar{T}}{4H_n}\right)^2}{\bar{T}}\right) \\ &\leq \exp\left(-\frac{\bar{T}}{4H_n^2}\right), \end{aligned}$$

where we used the lower bound on $\mathbb{E}[\bar{m}_i^+]$ shown in Lemma 3. Next, we derive with the help of Lemma 6 a statement of technical nature to be used later:

$$\begin{aligned} &\sum_{m=\lfloor \frac{\bar{T}}{4H_n} \rfloor + 1}^{\frac{\bar{T}}{2}} \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ = m\right) \\ &\leq 2 \sum_{m=\lfloor \frac{\bar{T}}{4H_n} \rfloor + 1}^{\frac{\bar{T}}{2}} \exp\left(-\frac{2m\varepsilon^2}{r_i^{+2}}\right) \\ &= 2 \sum_{m=0}^{\frac{\bar{T}}{2}} \exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right)^m - 2 \sum_{m=0}^{\lfloor \frac{\bar{T}}{4H_n} \rfloor} \exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right)^m \\ &= 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor} - \exp\left(-\frac{\varepsilon^2}{r_i^{+2}}\right)^{\bar{T}}}{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right) - 1} \\ &\leq 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right) - 1}. \end{aligned}$$

At last, putting both findings together, we derive our claim:

$$\begin{aligned} &\mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon\right) \\ &= \sum_{m=1}^{\frac{\bar{T}}{2}} \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ = m\right) \cdot \mathbb{P}(m_i^+ = m) \\ &= \sum_{m=1}^{\lfloor \frac{\bar{T}}{4H_n} \rfloor} \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ = m\right) \cdot \mathbb{P}(m_i^+ = m) + \sum_{m=\lfloor \frac{\bar{T}}{4H_n} \rfloor + 1}^{\frac{\bar{T}}{2}} \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ = m\right) \cdot \mathbb{P}(m_i^+ = m) \\ &\leq \mathbb{P}\left(\bar{m}_i^+ \leq \left\lfloor \frac{\bar{T}}{4H_n} \right\rfloor\right) + \sum_{m=\lfloor \frac{\bar{T}}{4H_n} \rfloor + 1}^{\frac{\bar{T}}{2}} \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon \mid m_i^+ = m\right) \\ &\leq \exp\left(-\frac{\bar{T}}{4H_n^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp\left(-\frac{2\varepsilon^2}{r_i^{+2}}\right) - 1}. \end{aligned}$$

□

Theorem 4 For any fixed player $i \in \mathcal{N}$ and $\varepsilon > 0$ the probability that the Shapley value estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater is bounded by

$$\mathbb{P}\left(|\hat{\phi}_i - \phi_i| \geq \varepsilon\right) \leq 2 \exp\left(-\frac{\bar{T}}{4H_n^2}\right) + 4 \frac{\exp\left(-\frac{2\varepsilon^2}{(r_i^+ + r_i^-)^2}\right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp\left(\frac{2\varepsilon^2}{(r_i^+ + r_i^-)^2}\right) - 1}.$$

Proof.

$$\begin{aligned} & \mathbb{P}\left(|\hat{\phi}_i - \phi_i| \geq \varepsilon\right) \\ &= \mathbb{P}\left(|(\hat{\phi}_i^+ - \phi_i^+) + (\hat{\phi}_i^- - \phi_i^-)| \geq \varepsilon\right) \\ &\leq \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| + |\hat{\phi}_i^- - \phi_i^-| \geq \varepsilon\right) \\ &\leq \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \frac{\varepsilon r_i^+}{r_i^+ + r_i^-}\right) + \mathbb{P}\left(|\hat{\phi}_i^- - \phi_i^-| \geq \frac{\varepsilon r_i^-}{r_i^+ + r_i^-}\right) \\ &\leq 2 \exp\left(-\frac{\bar{T}}{4H_n^2}\right) + 4 \frac{\exp\left(-\frac{2\varepsilon^2}{(r_i^+ + r_i^-)^2}\right)^{\lfloor \frac{\bar{T}}{4H_n} \rfloor}}{\exp\left(\frac{2\varepsilon^2}{(r_i^+ + r_i^-)^2}\right) - 1}. \end{aligned}$$

□

D Stratified SVARM Analysis

Notation:

- Let $\mathcal{L} = \{0, \dots, n-1\}$, $\mathcal{L}^+ = \{1, \dots, n-3\}$, and $\mathcal{L}^- = \{2, \dots, n-2\}$.
- Let $W = 2n + 1 + 2 \sum_{s=2}^{n-2} \lceil \frac{n}{s} \rceil$ denote the length of the warm-up phase.
- Let $\bar{T} = T - W$ be the available steps after the warm-up phase.
- Let $m_{i,\ell}^+ = \#\{t \mid i \in A_t, |A_t| = \ell + 1\}$ be the total number of samples used to update $\hat{\phi}_{i,\ell}^+$.
- Let $m_{i,\ell}^- = \#\{t \mid i \notin A_t, |A_t| = \ell\}$ be the total number of samples used to update $\hat{\phi}_{i,\ell}^-$.
- Let $\bar{m}_{i,\ell}^+ = \#\{t \mid i \in A_t, |A_t| = \ell + 1, t > W\}$ be the number of samples used to update $\hat{\phi}_{i,\ell}^+$ after the warm-up phase.
- Let $\bar{m}_{i,\ell}^- = \#\{t \mid i \notin A_t, |A_t| = \ell, t > W\}$ be the number of samples used to update $\hat{\phi}_{i,\ell}^-$ after the warm-up phase.
- Let $A_{i,\ell,k}^+$ be the k -th set used to update $\phi_{i,\ell}^+$ and $A_{i,\ell,k}^-$ the k -th set used to update $\phi_{i,\ell}^-$.
- Let $A_{i,\ell}$ be a random set with $\mathbb{P}(A_{i,\ell} = S) = \frac{1}{\binom{n-1}{\ell}}$ for all $S \subseteq \mathcal{N} \setminus \{i\}$ with $|S| = \ell$.
- Let $\hat{\phi}_{i,\ell}^+ = \frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \nu(A_{i,\ell,k}^+)$ and $\hat{\phi}_{i,\ell}^- = \frac{1}{m_{i,\ell}^-} \sum_{k=1}^{m_{i,\ell}^-} \nu(A_{i,\ell,k}^-)$.
- Let $\hat{\phi}_i = \frac{1}{n} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^+ - \hat{\phi}_{i,\ell}^-$.
- Let $\sigma_{i,\ell}^{+2} = \mathbb{V}[\nu(A_{i,\ell} \cup \{i\})]$ and $\sigma_{i,\ell}^{-2} = \mathbb{V}[\nu(A_{i,\ell})]$.
- Let $r_{i,\ell}^+ = \max_{S \subseteq \mathcal{N}: i \notin S, |S|=\ell} \nu(S \cup \{i\}) - \min_{S \subseteq \mathcal{N}: i \notin S, |S|=\ell} \nu(S \cup \{i\})$ be the range of $\nu(A_{i,\ell,k}^+)$.
- Let $r_{i,\ell}^- = \max_{S \subseteq \mathcal{N}: i \notin S, |S|=\ell} \nu(S) - \min_{S \subseteq \mathcal{N}: i \notin S, |S|=\ell} \nu(S)$ be the range of $\nu(A_{i,\ell,k}^-)$.
- Let $R_i^+ = \sum_{\ell=1}^{n-3} r_{i,\ell}^+$ and $R_i^- = \sum_{\ell=2}^{n-2} r_{i,\ell}^-$.

Assumptions:

- $n \geq 4$, for $n \leq 3$ the algorithm computes all Shapley values exactly.

D.1 Unbiasedness of Shapley Value Estimates

Lemma 8. *Due to the exact calculation, the following estimates are exact for all $i \in \mathcal{N}$:*

- $\hat{\phi}_{i,0}^+ = \phi_{i,0}^+ = \nu(\{i\})$
- $\hat{\phi}_{i,n-2}^+ = \phi_{i,n-2}^+ = \frac{1}{n-1} \sum_{j \in \mathcal{N}: j \neq i} \nu(\mathcal{N} \setminus \{j\})$
- $\hat{\phi}_{i,n-1}^+ = \phi_{i,n-1}^+ = \nu(\mathcal{N})$
- $\hat{\phi}_{i,0}^- = \phi_{i,0}^- = \nu(\emptyset) = 0$
- $\hat{\phi}_{i,1}^- = \phi_{i,1}^- = \frac{1}{n-1} \sum_{j \in \mathcal{N}: j \neq i} \nu(\{j\})$
- $\hat{\phi}_{i,n-1}^- = \phi_{i,n-1}^- = \nu(\mathcal{N} \setminus \{i\})$

Lemma 9. *All remaining estimates that are not calculated exactly are unbiased, i.e., for all $i \in \mathcal{N}$:*

- $\mathbb{E} [\hat{\phi}_{i,\ell}^+] = \phi_{i,\ell}^+$ for all $\ell \in \mathcal{L}^+$
- $\mathbb{E} [\hat{\phi}_{i,\ell}^-] = \phi_{i,\ell}^-$ for all $\ell \in \mathcal{L}^-$

Proof. We prove the statement only for $\hat{\phi}_{i,\ell}^+$ as the proof for $\hat{\phi}_{i,\ell}^-$ is analogous. Fix any $i \in \mathcal{N}$ and $\ell \in \mathcal{L}^+$. As soon as the size s_t of the to be sampled coalition $|A_t|$ is fixed, A_t is sampled uniformly from $\{S \subseteq \mathcal{N} \mid |S| = s_t\}$. This allows us to state for every A_t and any $S \subseteq \mathcal{N}$ with $|S| = \ell + 1$ and $i \notin S$:

$$\mathbb{P}(A_t = S \cup \{i\} \mid i \in A_t, |A_t| = \ell + 1) = \frac{1}{\binom{n-1}{\ell}}.$$

Continuing, we derive for the expectation of $\hat{\phi}_{i,\ell}^+$ given the number of samples $m_{i,\ell}^+$:

$$\begin{aligned} & \mathbb{E} [\hat{\phi}_{i,\ell}^+ \mid m_{i,\ell}^+] \\ &= \mathbb{E} \left[\frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \nu(A_{i,\ell,k}^+) \mid m_{i,\ell}^+ \right] \\ &= \frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \mathbb{E} [\nu(A_{i,\ell,k}^+) \mid m_{i,\ell}^+] \\ &= \frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \sum_{S \subseteq \mathcal{N} \setminus \{i\}; |S|=\ell} \mathbb{P}(A_{i,\ell,k}^+ = S \cup \{i\} \mid i \in A_{i,\ell,k}^+, |A_{i,\ell,k}^+| = \ell + 1) \cdot \nu(S \cup \{i\}) \\ &= \frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \sum_{S \subseteq \mathcal{N} \setminus \{i\}; |S|=\ell} \frac{1}{\binom{n-1}{\ell}} \cdot \nu(S \cup \{i\}) \\ &= \frac{1}{m_{i,\ell}^+} \sum_{k=1}^{m_{i,\ell}^+} \phi_{i,\ell}^+ \\ &= \phi_{i,\ell}^+. \end{aligned}$$

Note that the term is well defined, since $m_{i,\ell}^+ \in \{1, \dots, T\}$ due to the warm-up phase. We conclude:

$$\begin{aligned} \mathbb{E} [\hat{\phi}_{i,\ell}^+] &= \sum_{m=1}^T \mathbb{E} [\hat{\phi}_{i,\ell}^+ \mid m_{i,\ell}^+ = m] \cdot \mathbb{P}(m_{i,\ell}^+ = m) \\ &= \sum_{m=1}^T \phi_{i,\ell}^+ \cdot \mathbb{P}(m_{i,\ell}^+ = m) \\ &= \phi_{i,\ell}^+. \end{aligned}$$

□

Theorem 5 *The Shapley value estimates for all $i \in \mathcal{N}$ are unbiased, i.e.,*

$$\mathbb{E} [\hat{\phi}_i] = \phi_i.$$

Proof. By applying Lemma 8 and Lemma 9 we obtain:

$$\begin{aligned} \mathbb{E} [\hat{\phi}_i] &= \frac{1}{n} \sum_{\ell=0}^{n-1} \mathbb{E} [\hat{\phi}_{i,\ell}^+] - \mathbb{E} [\hat{\phi}_{i,\ell}^-] \\ &= \frac{1}{n} \sum_{\ell=0}^{n-1} \phi_{i,\ell}^+ - \phi_{i,\ell}^- \\ &= \phi_i. \end{aligned}$$

□

D.2 Sample numbers

Lemma 10. For any $i \in \mathcal{N}$ the number of updates $\bar{m}_{i,\ell}^+$ and $\bar{m}_{i,\ell}^-$ are binomially distributed with

$$\bar{m}_{i,\ell}^+ \sim \text{Bin}\left(\bar{T}, \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1)\right) \text{ for all } \ell \in \mathcal{L}^+ \\ \text{and } \bar{m}_{i,\ell}^- \sim \text{Bin}\left(\bar{T}, \frac{n-\ell}{n} \cdot \tilde{P}(\ell)\right) \text{ for all } \ell \in \mathcal{L}^-.$$

Proof. We argue that there are \bar{T} many independent time steps in which $\hat{\phi}_{i,\ell}^+$ can be updated. If $|A_t| = \ell + 1$ then i is included in A_t with a probability of $\frac{\ell+1}{n}$ due to the uniform sampling of A_t given that its size s_t is fixed, leading to an update. Since the choice of size and members of the set A_t are independent, the probability of $\hat{\phi}_{i,\ell}^+$ being updated in time step t is $\frac{\ell+1}{n} \cdot \tilde{P}(\ell+1)$. The same argument holds true for $\hat{\phi}_{i,\ell}^-$ with an update probability of $\frac{n-\ell}{n} \cdot \tilde{P}(\ell)$ in each time step. □

Lemma 11. For any $i \in \mathcal{N}$ the expected number of updates of $\hat{\phi}_{i,\ell}^+$ and $\hat{\phi}_{i,\ell}^-$ after the warm-up phase is at least

$$\mathbb{E}[\bar{m}_{i,\ell}^+] \geq \frac{\bar{T}}{2n \log n} \text{ for all } \ell \in \mathcal{L}^+ \\ \text{and } \mathbb{E}[\bar{m}_{i,\ell}^-] \geq \frac{\bar{T}}{2n \log n} \text{ for all } \ell \in \mathcal{L}^-.$$

Proof. In the following we distinguish between different cases, depending on the parity of n and size of ℓ . We will use the bound $H_n \leq \log n + 1$ and the following inequalities multiple times which hold true for $n \geq 4$:

$$\frac{1}{H_{\frac{n-1}{2}} - 1} \geq \frac{1}{\log n} \quad \text{and} \quad \frac{n \log n - 1}{n(H_{\frac{n}{2}-1} - 1)} \geq 1.$$

We begin with the case of $n \nmid 2$ and $\ell \leq \frac{n-1}{2} - 1$:

$$\begin{aligned} \mathbb{E}[\bar{m}_{i,\ell}^+] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) \\ &= \frac{\bar{T}}{2n} \cdot \frac{1}{H_{\frac{n-1}{2}} - 1} \\ &\geq \frac{\bar{T}}{2n \log n} \end{aligned} \quad \begin{aligned} \mathbb{E}[\bar{m}_{i,\ell}^-] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n} \cdot \frac{n-\ell}{\ell} \cdot \frac{1}{H_{\frac{n-1}{2}} - 1} \\ &\geq \frac{\bar{T}}{2n \log n} \end{aligned}$$

For $n \nmid 2$ and $\ell = \frac{n-1}{2}$ we obtain:

$$\begin{aligned} \mathbb{E}[\bar{m}_{i,\ell}^+] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) \\ &= \frac{\bar{T}}{2n} \cdot \frac{\ell+1}{n-\ell-1} \cdot \frac{1}{H_{\frac{n-1}{2}} - 1} \\ &\geq \frac{\bar{T}}{2n \log n} \end{aligned} \quad \begin{aligned} \mathbb{E}[\bar{m}_{i,\ell}^-] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n} \cdot \frac{n-\ell}{\ell} \cdot \frac{1}{H_{\frac{n-1}{2}} - 1} \\ &\geq \frac{\bar{T}}{2n \log n}. \end{aligned}$$

For $n \nmid 2$ and $\ell \geq \frac{n+1}{2}$ we obtain:

$$\begin{aligned} \mathbb{E} \left[\bar{m}_{i,\ell}^+ \right] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) & \mathbb{E} \left[\bar{m}_{i,\ell}^- \right] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n} \cdot \frac{\ell+1}{n-\ell-1} \cdot \frac{1}{H_{\frac{n-1}{2}}-1} & &= \frac{\bar{T}}{2n} \cdot \frac{1}{H_{\frac{n-1}{2}}-1} \\ &\geq \frac{\bar{T}}{2n \log n} & &\geq \frac{\bar{T}}{2n \log n} \end{aligned}$$

Switching to $n \mid 2$, we start with $\ell = \frac{n}{2} - 1$:

$$\begin{aligned} \mathbb{E} \left[\bar{m}_{i,\ell}^+ \right] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) & \mathbb{E} \left[\bar{m}_{i,\ell}^- \right] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n \log n} & &= \frac{\bar{T}}{2n \log n} \cdot \frac{n-\ell}{\ell} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} \\ & & &\geq \frac{\bar{T}}{2n \log n} \end{aligned}$$

For $n \mid 2$ and $\ell = \frac{n}{2}$ we derive:

$$\begin{aligned} \mathbb{E} \left[\bar{m}_{i,\ell}^+ \right] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) & \mathbb{E} \left[\bar{m}_{i,\ell}^- \right] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n \log n} \cdot \frac{\ell+1}{n-\ell-1} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} & &= \frac{\bar{T}}{2n \log n} \\ &\geq \frac{\bar{T}}{2n \log n} & & \end{aligned}$$

For $n \mid 2$ and $\ell \leq \frac{n}{2} - 2$ we derive:

$$\begin{aligned} \mathbb{E} \left[\bar{m}_{i,\ell}^+ \right] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) & \mathbb{E} \left[\bar{m}_{i,\ell}^- \right] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n \log n} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} & &= \frac{\bar{T}}{2n \log n} \cdot \frac{n-\ell}{\ell} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} \\ &\geq \frac{\bar{T}}{2n \log n} & &\geq \frac{\bar{T}}{2n \log n} \end{aligned}$$

Finally, $n \mid 2$ and $\ell \geq \frac{n}{2} + 1$ yields:

$$\begin{aligned} \mathbb{E} \left[\bar{m}_{i,\ell}^+ \right] &= \bar{T} \cdot \frac{\ell+1}{n} \cdot \tilde{P}(\ell+1) & \mathbb{E} \left[\bar{m}_{i,\ell}^- \right] &= \bar{T} \cdot \frac{n-\ell}{n} \cdot \tilde{P}(\ell) \\ &= \frac{\bar{T}}{2n \log n} \cdot \frac{\ell+1}{n-\ell-1} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} & &= \frac{\bar{T}}{2n \log n} \cdot \frac{n \log n - 1}{n(H_{\frac{n}{2}-1}-1)} \\ &\geq \frac{\bar{T}}{2n \log n} & &\geq \frac{\bar{T}}{2n \log n} \end{aligned}$$

□

D.3 Variance and Expected Squared Error

Lemma 12. *The variance of any player's Shapley value estimate $\hat{\phi}_i$ given the number of samples $m_{i,\ell}^+$ and $m_{i,\ell}^-$ for all $\ell \in \mathcal{L}$ is given by*

$$\mathbb{V} \left[\hat{\phi}_i \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+}, \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] = \frac{1}{n^2} \sum_{\ell=1}^{n-3} \frac{\sigma_{i,\ell}^{+2}}{m_{i,\ell}^+} + \frac{\sigma_{i,\ell+1}^{-2}}{m_{i,\ell+1}^-}.$$

Proof. We first decompose the variance of $\hat{\phi}_i$ into the variances of $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ and their covariance:

$$\begin{aligned} & \mathbb{V} \left[\hat{\phi}_i \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+}, \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] \\ &= \mathbb{V} \left[\hat{\phi}_i^+ \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+} \right] + \mathbb{V} \left[\hat{\phi}_i^- \mid \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] - 2\text{Cov} \left(\hat{\phi}_i^+, \hat{\phi}_i^- \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+}, \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right) \\ &= \mathbb{V} \left[\hat{\phi}_i^+ \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+} \right] + \mathbb{V} \left[\hat{\phi}_i^- \mid \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right]. \end{aligned}$$

where we used the observation that $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ are independent. We derive for $\hat{\phi}_i^+$:

$$\begin{aligned} \mathbb{V} \left[\hat{\phi}_i^+ \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+} \right] &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \mathbb{V} \left[\hat{\phi}_{i,\ell}^+ \mid m_{i,\ell}^+ \right] + \sum_{\ell \neq \ell'} \text{Cov} \left(\hat{\phi}_{i,\ell}^+, \hat{\phi}_{i,\ell'}^+ \mid m_{i,\ell}^+, m_{i,\ell'}^+ \right) \\ &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \mathbb{V} \left[\hat{\phi}_{i,\ell}^+ \mid m_{i,\ell}^+ \right] \\ &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \frac{\sigma_{i,\ell}^{+2}}{m_{i,\ell}^+}, \end{aligned}$$

where we used the observation that $\hat{\phi}_{i,\ell}^+$ and $\hat{\phi}_{i,\ell'}^+$ are independent for $\ell \neq \ell'$. Note that $\hat{\phi}_{i,0}^+$, $\hat{\phi}_{i,n-2}^+$, $\hat{\phi}_{i,n-1}^+$, $\hat{\phi}_{i,0}^-$, $\hat{\phi}_{i,1}^-$, and $\hat{\phi}_{i,n-1}^-$ are constants without variance. A similar result can be obtained for $\hat{\phi}_i^-$. Putting our intermediate results together yields:

$$\begin{aligned} & \mathbb{V} \left[\hat{\phi}_i \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+}, \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] \\ &= \mathbb{V} \left[\hat{\phi}_i^+ \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+} \right] + \mathbb{V} \left[\hat{\phi}_i^- \mid \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] \\ &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \frac{\sigma_{i,\ell}^{+2}}{m_{i,\ell}^+} + \frac{1}{n^2} \sum_{\ell=2}^{n-2} \frac{\sigma_{i,\ell}^{-2}}{m_{i,\ell}^-} \\ &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \frac{\sigma_{i,\ell}^{+2}}{m_{i,\ell}^+} + \frac{\sigma_{i,\ell+1}^{-2}}{m_{i,\ell+1}^-}. \end{aligned}$$

□

Lemma 13. For any $i \in \mathcal{N}$ holds

$$\mathbb{E} \left[\frac{1}{m_{i,\ell}^+} \right] \leq \frac{2n \log n}{\bar{T}} \text{ for all } \ell \in \mathcal{L}^+ \quad \text{and} \quad \mathbb{E} \left[\frac{1}{m_{i,\ell}^-} \right] \leq \frac{2n \log n}{\bar{T}} \text{ for all } \ell \in \mathcal{L}^-.$$

Proof. We prove the result only for $m_{i,\ell}^+$ since the proof for $m_{i,\ell}^-$ is analogous. By combining Equation (3.4) in Chao and Strawderman [1972]:

$$\mathbb{E} \left[\frac{1}{1+X} \right] = \frac{1 - (1-p)^{m+1}}{(m+1)p} \leq \frac{1}{mp} = \frac{1}{\mathbb{E}[X]},$$

for any binomially distributed random variable $X \sim \text{Bin}(m, p)$ with Lemma 10 and Lemma 11, we obtain:

$$\mathbb{E} \left[\frac{1}{m_{i,\ell}^+} \right] = \mathbb{E} \left[\frac{1}{1 + \bar{m}_{i,\ell}^+} \right] \leq \frac{1}{\mathbb{E}[\bar{m}_{i,\ell}^+]} \leq \frac{2n \log n}{\bar{T}}.$$

□

Theorem 6 For \tilde{P} as chosen above the variance of any player's Shapley value estimate $\hat{\phi}_i$ is bounded by

$$\mathbb{V} \left[\hat{\phi}_i \right] \leq \frac{2 \log n}{n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

Proof. The combination of Lemma 12 and Lemma 13 yields:

$$\begin{aligned}
 \mathbb{V}[\hat{\phi}_i] &= \mathbb{E} \left[\mathbb{V} \left[\hat{\phi}_i \mid \left(m_{i,\ell}^+ \right)_{\ell \in \mathcal{L}^+}, \left(m_{i,\ell}^- \right)_{\ell \in \mathcal{L}^-} \right] \right] \\
 &\leq \mathbb{E} \left[\frac{1}{n^2} \sum_{\ell=1}^{n-3} \frac{\sigma_{i,\ell}^{+2}}{m_{i,\ell}^+} + \frac{\sigma_{i,\ell+1}^{-2}}{m_{i,\ell+1}^-} \right] \\
 &= \frac{1}{n^2} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} \cdot \mathbb{E} \left[\frac{1}{m_{i,\ell}^+} \right] + \sigma_{i,\ell+1}^{-2} \cdot \mathbb{E} \left[\frac{1}{m_{i,\ell+1}^-} \right] \\
 &\leq \frac{2 \log n}{n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.
 \end{aligned}$$

□

Corollary 2 For \tilde{P} as chosen above the MSE of any player's Shapley value estimate $\hat{\phi}_i$ is bounded by

$$\mathbb{E} \left[\left(\hat{\phi}_i - \phi_i \right)^2 \right] \leq \frac{2 \log n}{n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

Proof. Using the bias-variance decomposition, the unbiasedness of $\hat{\phi}_i$ shown in Theorem 5, and the bound on the variance from Theorem 6 we obtain that:

$$\begin{aligned}
 \mathbb{E} \left[\left(\hat{\phi}_i - \phi_i \right)^2 \right] &= \left(\mathbb{E}[\hat{\phi}_i] - \phi_i \right)^2 + \mathbb{V}[\hat{\phi}_i] \\
 &\leq \frac{2 \log n}{n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.
 \end{aligned}$$

□

D.4 Probabilistic Bounds

Theorem 7 Fix any player $i \in \mathcal{N}$ and $\varepsilon > 0$. For \tilde{P} as above the probability that the estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater is bounded by

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \frac{2 \log n}{\varepsilon^2 n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

Proof. The bound on the variance of $\hat{\phi}_i$ in Theorem 6 allows us to apply Chebyshev's inequality:

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \frac{\mathbb{V}[\hat{\phi}_i]}{\varepsilon^2} \leq \frac{2 \log n}{\varepsilon^2 n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}.$$

□

Corollary 4. Fix any player $i \in \mathcal{N}$ and $\delta \in (0, 1]$. The estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater with probability not greater than δ , i.e.,

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \varepsilon \right) \leq \delta \quad \text{for} \quad \varepsilon = \sqrt{\frac{2 \log n}{\delta n\bar{T}} \sum_{\ell=1}^{n-3} \sigma_{i,\ell}^{+2} + \sigma_{i,\ell+1}^{-2}}.$$

Lemma 14. For any $i \in \mathcal{N}$ and fixed $\varepsilon > 0$ holds:

$$\bullet \mathbb{P}(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+) \leq 2 \exp \left(-\frac{2m_{i,\ell}^+ \varepsilon^2}{r_{i,\ell}^{+2}} \right) \text{ for all } \ell \in \mathcal{L}^+$$

$$\bullet \mathbb{P}(|\hat{\phi}_{i,\ell}^- - \phi_{i,\ell}^-| \geq \varepsilon \mid m_{i,\ell}^-) \leq 2 \exp\left(-\frac{2m_{i,\ell}^- \varepsilon^2}{r_{i,\ell}^-^2}\right) \text{ for all } \ell \in \mathcal{L}^-$$

Proof. We prove the statement for $\hat{\phi}_{i,\ell}^+$ by making use of Hoeffding's inequality in combination with the unbiasedness of the strata estimates shown in Lemma 9. The proof for $\hat{\phi}_{i,\ell}^-$ is analogous.

$$\begin{aligned} & \mathbb{P}(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+) \\ &= \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \mathbb{E}[\hat{\phi}_{i,\ell}^+]| \geq \varepsilon \mid m_{i,\ell}^+\right) \\ &= \mathbb{P}\left(\left|\sum_{k=1}^{m_{i,\ell}^+} \nu(A_{i,\ell,k}^+) - \mathbb{E}\left[\sum_{k=1}^{m_{i,\ell}^+} \nu(A_{i,\ell,k}^+)\right]\right| \geq m_{i,\ell}^+ \varepsilon \mid m_{i,\ell}^+\right) \\ &\leq 2 \exp\left(-\frac{2m_{i,\ell}^+ \varepsilon^2}{r_{i,\ell}^+^2}\right). \end{aligned}$$

□

Lemma 15. For any $i \in \mathcal{N}$ and fixed $\varepsilon > 0$ holds:

$$\begin{aligned} \bullet \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon\right) &\leq \exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^+^2}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2}{r_{i,\ell}^+^2}\right) - 1} \text{ for all } \ell \in \mathcal{L}^+ \\ \bullet \mathbb{P}\left(|\hat{\phi}_{i,\ell}^- - \phi_{i,\ell}^-| \geq \varepsilon\right) &\leq \exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^-^2}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2}{r_{i,\ell}^-^2}\right) - 1} \text{ for all } \ell \in \mathcal{L}^- \end{aligned}$$

Proof. We prove the statement for $\hat{\phi}_{i,\ell}^+$. The proof for $\hat{\phi}_{i,\ell}^-$ is analogous. To begin with, we derive with the help of Hoeffding's inequality for binomial distributions a bound for the probability of $\bar{m}_{i,\ell}^+$ not exceeding $\frac{\bar{T}}{4n \log n}$:

$$\begin{aligned} & \mathbb{P}\left(\bar{m}_{i,\ell}^+ \leq \frac{\bar{T}}{4n \log n}\right) \\ &\leq \mathbb{P}\left(\mathbb{E}[\bar{m}_{i,\ell}^+] - \bar{m}_{i,\ell}^+ \geq \mathbb{E}[\bar{m}_{i,\ell}^+] - \frac{\bar{T}}{4n \log n}\right) \\ &\leq \exp\left(-\frac{2\left(\mathbb{E}[\bar{m}_{i,\ell}^+] - \frac{\bar{T}}{4n \log n}\right)^2}{\bar{T}}\right) \\ &\leq \exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right), \end{aligned}$$

where we used the lower bound on $\mathbb{E}[\bar{m}_{i,\ell}^+]$ shown in Lemma 11. Next, we derive with the help of Lemma 14 a statement of technical nature to be used later:

$$\begin{aligned}
 & \sum_{m=\lfloor \frac{\bar{T}}{4n \log n} \rfloor + 1}^{\bar{T}} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+ = m\right) \\
 & \leq 2 \sum_{m=\lfloor \frac{\bar{T}}{4n \log n} \rfloor + 1}^{\bar{T}} \exp\left(-\frac{2m\varepsilon^2}{r_{i,\ell}^{+2}}\right) \\
 & = 2 \sum_{m=0}^{\bar{T}} \exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^m - 2 \sum_{m=0}^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor} \exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^m \\
 & = 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor} - \exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^{\bar{T}}}{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right) - 1} \\
 & \leq 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right) - 1}.
 \end{aligned}$$

At last, putting both findings together, we derive our claim:

$$\begin{aligned}
 & \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon\right) \\
 & = \sum_{m=1}^{\bar{T}} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+ = m\right) \cdot \mathbb{P}\left(m_{i,\ell}^+ = m\right) \\
 & = \sum_{m=1}^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+ = m\right) \cdot \mathbb{P}\left(m_{i,\ell}^+ = m\right) \\
 & \quad + \sum_{m=\lfloor \frac{\bar{T}}{4n \log n} \rfloor + 1}^{\bar{T}} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+ = m\right) \cdot \mathbb{P}\left(m_{i,\ell}^+ = m\right) \\
 & \leq \mathbb{P}\left(\bar{m}_{i,\ell}^+ \leq \left\lfloor \frac{\bar{T}}{4n \log n} \right\rfloor\right) + \sum_{m=\lfloor \frac{\bar{T}}{4n \log n} \rfloor + 1}^{\bar{T}} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon \mid m_{i,\ell}^+ = m\right) \\
 & \leq \exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(-\frac{2\varepsilon^2}{r_{i,\ell}^{+2}}\right) - 1}.
 \end{aligned}$$

□

Lemma 16. For any $i \in \mathcal{N}$ and fixed $\varepsilon > 0$ the probabilities that the estimates $\hat{\phi}_i^+$ and $\hat{\phi}_i^-$ deviate from ϕ_i^+ , respectively ϕ_i^- are bounded by:

$$\bullet \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon\right) \leq (n-3) \left(\exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2 n^2}{R_i^{+2}}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(-\frac{2\varepsilon^2 n^2}{R_i^{+2}}\right) - 1} \right),$$

$$\bullet \mathbb{P}\left(|\hat{\phi}_i^- - \phi_i^-| \geq \varepsilon\right) \leq (n-3) \left(\exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2 n^2}{R_i^-}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2 n^2}{R_i^-}\right) - 1} \right).$$

Proof. We prove the statement for $\hat{\phi}_i^+$ using Lemma 15. The proof for $\hat{\phi}_i^-$ is analogous.

$$\begin{aligned} & \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \varepsilon\right) \\ &= \mathbb{P}\left(\left|\frac{1}{n} \sum_{\ell=0}^{n-1} \hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+\right| \geq \varepsilon\right) \\ &\leq \mathbb{P}\left(\frac{1}{n} \sum_{\ell=0}^{n-1} |\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon\right) \\ &= \mathbb{P}\left(\sum_{\ell=1}^{n-3} |\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \varepsilon n\right) \\ &\leq \sum_{\ell=1}^{n-3} \mathbb{P}\left(|\hat{\phi}_{i,\ell}^+ - \phi_{i,\ell}^+| \geq \frac{\varepsilon n r_{i,\ell}^+}{R^+}\right) \\ &\leq (n-3) \left(\exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2 n^2}{R^+}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2 n^2}{R^+}\right) - 1} \right) \end{aligned}$$

□

Theorem 8 For any $i \in \mathcal{N}$ and fixed $\varepsilon > 0$ the probability that the estimate $\hat{\phi}_i$ deviates from ϕ_i by a margin of ε or greater is bounded by

$$\mathbb{P}\left(|\hat{\phi}_i - \phi_i| \geq \varepsilon\right) \leq 2(n-3) \left(\exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2 n^2}{(R_i^+ + R_i^-)^2}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2 n^2}{(R_i^+ + R_i^-)^2}\right) - 1} \right).$$

Proof. We apply Lemma 16 and obtain:

$$\begin{aligned} & \mathbb{P}\left(|\hat{\phi}_i - \phi_i| \geq \varepsilon\right) \\ &= \mathbb{P}\left(|(\hat{\phi}_i^+ - \phi_i^+) + (\hat{\phi}_i^- - \phi_i^-)| \geq \varepsilon\right) \\ &\leq \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| + |\hat{\phi}_i^- - \phi_i^-| \geq \varepsilon\right) \\ &\leq \mathbb{P}\left(|\hat{\phi}_i^+ - \phi_i^+| \geq \frac{\varepsilon R_i^+}{R_i^+ + R_i^-}\right) + \mathbb{P}\left(|\hat{\phi}_i^- - \phi_i^-| \geq \frac{\varepsilon R_i^-}{R_i^+ + R_i^-}\right) \\ &\leq 2(n-3) \left(\exp\left(-\frac{\bar{T}}{8n^2(\log n)^2}\right) + 2 \frac{\exp\left(-\frac{2\varepsilon^2 n^2}{(R_i^+ + R_i^-)^2}\right)^{\lfloor \frac{\bar{T}}{4n \log n} \rfloor}}{\exp\left(\frac{2\varepsilon^2 n^2}{(R_i^+ + R_i^-)^2}\right) - 1} \right). \end{aligned}$$

□

E Cooperative Games

E.1 Synthetic games

We provide formal definitions of the synthetic games and their Shapley values used in our empirical evaluation (see Section 6), and describe the process of how we randomly generated some of these.

E.1.1 Shoe Game

The number of players n in the Shoe game has to be even. The player set consist of two halves A and B of equal size, i.e., $\mathcal{N} = A \cup B$ with $A \cap B = \emptyset$ and $|A| = |B| = \frac{n}{2}$. The value function is given by $\nu(S) = \min\{|S \cap A|, |S \cap B|\}$. All players share the same Shapley value of $\phi_i = \frac{1}{2}$.

E.1.2 Airport Game

The Airport game entails $n = 100$ players. Each player i has an assigned weight c_i . The value function is the maximum of all weights contained in the coalition, i.e., $\nu(S) = \max_{i \in S} c_i$. The weights and resulting Shapley values are:

$$c_i = \begin{cases} 1 & \text{if } i \in \{1, \dots, 8\} \\ 2 & \text{if } i \in \{9, \dots, 20\} \\ 3 & \text{if } i \in \{21, \dots, 26\} \\ 4 & \text{if } i \in \{27, \dots, 40\} \\ 5 & \text{if } i \in \{41, \dots, 48\} \\ 6 & \text{if } i \in \{49, \dots, 57\} \\ 7 & \text{if } i \in \{58, \dots, 70\} \\ 8 & \text{if } i \in \{71, \dots, 80\} \\ 9 & \text{if } i \in \{81, \dots, 90\} \\ 10 & \text{if } i \in \{91, \dots, 100\} \end{cases}, \quad \phi_i = \begin{cases} 0.01 & \text{if } i \in \{1, \dots, 8\} \\ 0.020869565 & \text{if } i \in \{9, \dots, 20\} \\ 0.033369565 & \text{if } i \in \{21, \dots, 26\} \\ 0.046883079 & \text{if } i \in \{27, \dots, 40\} \\ 0.063549745 & \text{if } i \in \{41, \dots, 48\} \\ 0.082780515 & \text{if } i \in \{49, \dots, 57\} \\ 0.106036329 & \text{if } i \in \{58, \dots, 70\} \\ 0.139369662 & \text{if } i \in \{71, \dots, 80\} \\ 0.189369662 & \text{if } i \in \{81, \dots, 90\} \\ 0.289369662 & \text{if } i \in \{91, \dots, 100\} \end{cases}.$$

E.1.3 SOUG Game

A Sum of unanimity games (SOUG) is specified by M many sets $S_1, \dots, S_M \subseteq \mathcal{N}$ and weights $c_1, \dots, c_M \in \mathbb{R}$. The value functions is defined as $\nu(S) = \sum_{m=1}^M c_m \cdot \mathbb{I}_{S_m \subseteq S}$ leading to Shapley values $\phi_i = \sum_{m=1}^M \frac{c_m}{|S_m|} \cdot \mathbb{I}_{i \in S_m}$, which can be computed in polynomial time if knowledge of sets and coefficients is provided. We generate SOUG games with $M = 50$ randomly by selecting for each S_m to be drawn a size uniformly at random between 1 and n , and then draw the set S_m with that size uniformly. We draw the coefficients uniformly at random from $[0, 1]$.

E.2 Explainability games

In the following, we describe the three explainability games introduced in Section 6; namely, the NLP sentiment analysis game (see Section E.2.1), the image classifier game (see Section E.2.2), and the adult classification (see Section E.2.3), and explain the value function of each resulting cooperative game. Since there exists no efficient closed-form solution, we compute the Shapley values exhaustively (via brute force) in order to allow the tracking of the approximation error of the different algorithms. Due to constraints in computational power this limits us to $n = 14$ players per game, which necessitates $2^{14} = 16\,384$ model evaluation to exhaustively traverse the powerset of all coalitions. Note that after a budget of $T = 2^{14}$ both KernelSHAP and Stratified SVARM⁺ have an approximation error of zero because both have observed all coalition values and thus have seen the cooperative game in its entirety.

E.2.1 NLP sentiment analysis

The NLP sentiment game describes an explainability scenario for local feature importance of a sentiment classification model. The sentiment classifier³ is a fine-tuned version of the DistilBERT transformer architecture Sanh et al. [2019]. The model was fine-tuned on the *IMDB* dataset Maas et al. [2011]. The model expects a natural language sentence as input, transforms the sentence into a tokenized form and predicts a sentiment score ranging from $[-1, 1]$. We randomly select sentences from the *IMDB* dataset, that contain no more than 14 tokens. For a sentence the local explainability

³<https://huggingface.co/dhlee347/distilbert-imdb>.

game consists of presenting the model a coalition of players (tokens) and observing the predicted sentiment as the value of a coalition. Absent players are removed in the tokenized representation (i.e. tokens are removed).

E.2.2 Image classifier

The image classifier game is similar to the NLP sentiment analysis game (Section E.2.1) a local explanation scenario. For this we explain the output of an image classifier given random images from ImageNet Deng et al. [2009]. The model to be explained is a ResNet18⁴ He et al. [2016] trained on ImageNet Deng et al. [2009]. To restrict the number of players, we apply SLIC Achanta et al. [2012] to summarize individual pixels with 14 super-pixels. The super-pixels then make up the players in the image classification game. A coalition of players, thus, consists of the corresponding super-pixels. The super-pixels of absent players are removed via mean-imputation by setting all their pixels to grey. The worth of a coalition is determined by the output of the model (using only the present super-pixels given by the coalition) for the class of the original prediction which was made with all pixels being present.

E.2.3 Adult classification

Similar to the preceding two games, the adult classification game is also a local explanation scenario. We train a gradient-boosted tree classifier (sklearn) on the adult dataset Becker and Kohavi [1996] to classify whether an adult has an income below or above 50 000. Each game is based on a randomly chosen datapoint for which the players correspond to features. A coalition is formed by removing the absent feature values of the selected datapoint via mean imputation. The worth of a coalition is the predicted class probability of the true income class given the manipulated datapoint after mean imputation of absent features.

⁴<https://pytorch.org/vision/main/models/generated/torchvision.models.resnet18.html>.

F Further Empirical Results

The plots shown in Figure 3 hardly visualize the performance differences between the algorithms with low MSE values. Thus, we present our findings in Figure 4 to Figure 9 in higher resolution.




Figure 4: Airport game with 100 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.




Figure 5: SOUG game with 20 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.




Figure 6: Shoe game with 50 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.




Figure 7: NLP game with 14 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.




Figure 8: Image classifier game with 14 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.

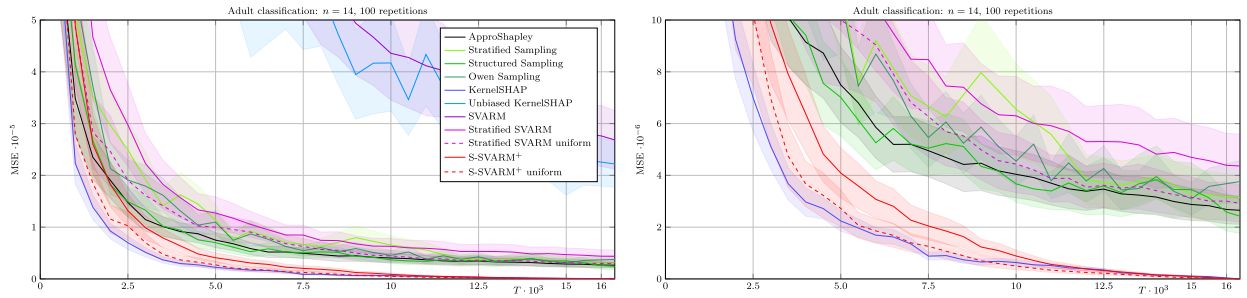


Figure 9: Adult classification with 14 players: Averaged MSE over 100 repetitions in dependence of fixed budget T , shaded bands showing standard errors.