

Identifying Top-k Players in Cooperative Games via Shapley Bandits

Patrick Kolpaczki¹, Viktor Bengs¹ and Eyke Hüllermeier²

¹Paderborn University, Germany

²University of Munich (LMU), Germany

Abstract

The usefulness of cooperative game theory and key concepts like the Shapley value, which measures the contribution of individual players to the overall performance of a coalition, has been demonstrated in various applications. Due to the computational effort growing exponentially with the number of participants in a game, several methods have been proposed to approximate Shapley values. Yet, in many applications, only the order of players according to their Shapley values is important, or maybe the set of the k best players, but not the values themselves. In this paper, we consider the problem of identifying the k players in a cooperative game with the highest Shapley values and denote it as the *Top-k Shapley* problem. By viewing the marginal contributions of a player as a random variable, we establish a connection between cooperative games and multi-armed bandits, which in turn allows us to reduce Top- k Shapley to the multiple arms identification problem. We call the resulting bandits problem *Shapley bandits*. Besides adopting existing algorithms for multiple arms identifications, we propose the *Border Uncertainty Sampling* algorithm (BUS) and provide empirical evidence for its superiority over state-of-the-art algorithms.

Keywords

Shapley value, Cooperative games, Multi-armed bandit, Multiple arms identification

1. Introduction

The formal notion of a cooperative game, in which players can form coalitions to accomplish a certain task, is a versatile concept with countless practical applications. Consider, for example, the cooperation of municipalities in infrastructure projects, with the goal to reduce costs by sharing and allocating available resources. In the context of (supervised) machine learning, individual features can be seen as players and feature subsets as coalitions – the task here is to train a model with high predictive performance [1, 2].

An interesting question in the context of cooperative games concerns the importance or contribution of an individual player: How to distribute the collective benefit of a coalition among the individual players? A connection to explainable AI can be drawn by interpreting features in a machine learning model as players and the predictive performance as the collective benefit such that the portion allocated to each feature can be seen as its importance for the model. Independent of the considered application, cooperative game theory has proposed different solution concepts, with the Shapley value as the arguably most popular one [3]. The Shapley

LWDA'21: Lernen, Wissen, Daten, Analysen September 01–03, 2021, Munich, Germany

✉ patrick.kolpaczki@upb.de (P. Kolpaczki); viktor.bengs@upb.de (V. Bengs); eyke@ifi.lmu.de (E. Hüllermeier)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

value assigns to each player a weighted average of all its marginal contributions, where we understand by a marginal contribution of a player the increase in the worth of a coalition when adding that player. The popularity of the Shapley value arises from the fact that it can be derived axiomatically by demanding desirable properties that one would expect from a fair distribution [3]. It has found its usage in a broad range of fields, from identifying influential members in terrorist networks [4, 5] to finding important neurons in artificial neural networks [6].

An inherent drawback of the Shapley value is the huge computational effort caused by the exponentially (in the number of players) growing number of marginal contributions — one per coalition — to be averaged over. As a consequence, brute force approaches quickly become infeasible for even only a few dozens of players. Several approximation methods have been proposed [6, 7, 8] to tackle this difficulty, all of them sharing the same idea of calculating mean estimates for randomly sampled marginal contributions uniformly for all players. Further, theoretical guarantees for approximation methods have been shown under mild assumptions [7, 8].

While these approximations show partially satisfying results in empirical studies, it seems to be rarely mentioned that in many applications the true objective is not to obtain precise Shapley value estimates for all players, but to identify a certain number of k players with the highest Shapley values (even though most works are indirectly aiming for that). For example, security agencies are more interested in identifying the most threatening members in terrorist networks, or the good performance of a machine learning model is oftentimes largely driven only by the most valuable features. Needless to say, one could tackle this problem naïvely by just pointing at the k players with highest Shapley value estimates obtained by traditional approximation algorithms. However, this approach would involve sampling steps to approximate the Shapley values of players for which one can already be certain that these are at the top or bottom of the ranking in terms of the Shapley values. In such cases, on the other hand, it makes sense to sample marginal contributions for players lying in the "middle" of the ranking in order to separate as quickly as possible the set of k -best players from the rest with a certain degree of certainty, although this might involve sacrificing precision of estimates for those players who are likely to be at the top or bottom of the ranking.

Similar considerations have already been made in the field of multi-armed bandit (MAB) problems [9], which is a class of online learning problems, where an agent needs to choose one arm (choice alternative) among a given set of arms (choice alternatives) in the course of a sequential decision process to achieve a specific target. In the stochastic variant of the MAB problem, each arm is associated with an unknown reward distribution and choosing a specific arm results in obtaining a stochastic reward generated by the chosen arm's unknown reward distribution. Many of the targets considered therefore revolve around identifying a specific partial ranking with respect to the (unknown) means of the arms reward distributions as quickly as possible. One particular target is to find the k arms having the highest mean, known as the multiple arms identifications problem [10], for which a number of algorithmic solutions are already available [10, 11, 12, 13, 14]. In this paper, we show how to trace the Top- k Shapley problem back to the multiple arms identifications problem, so that state-of-the-art solution methods for the latter problem can be efficiently used for the former. In addition, we propose a new method that performs even superior in numerical experiments.

2. Preliminaries

Before introducing our proposed problem formally in Section 3, we revisit in the following cooperative games and the Shapley value, as well as the problem of multiple arms identification in multi-armed bandit problems.

2.1. Cooperative Games and the Shapley Value

A *cooperative game* is characterized by a pair (N, ν) containing a set of *players* $N = \{p_1, \dots, p_n\}$ and a *value function* $\nu : \mathcal{P}(N) \rightarrow \mathbb{R}$, where $\nu(\emptyset) = 0$ by definition. The players can form *coalitions* $S \subseteq N$ and obtain a combined benefit given by $\nu(S)$ which is called the *worth* of S . For the question of how to distribute the worth $\nu(N)$ of the *grand coalition* N to the individual n many players, the *Shapley value* [3] forms a payoff distribution allocating to each player p_i the value

$$\phi_i(\nu) = \sum_{S \subseteq N \setminus \{p_i\}} \frac{1}{n \binom{n-1}{|S|}} \cdot (\nu(S \cup \{p_i\}) - \nu(S)).$$

For simplicity, we write ϕ_i whenever it is clear to which value function ν we refer. The difference in worth $\nu(S \cup \{p_i\}) - \nu(S)$ is called p_i 's *marginal contribution* given S . The Shapley value can be derived axiomatically, as it is provably the only solution concept fulfilling simultaneously the following properties [3], which one would intuitively demand from a fair distribution:

- **Efficiency:** the worth of N is partitioned over all players, i.e., $\nu(N) = \sum_{p_i \in N} \phi_i$,
- **Symmetry:** if two players p_i and p_j cannot be distinguished by their marginal contributions, i.e., $\nu(S \cup \{p_i\}) = \nu(S \cup \{p_j\})$ for all $S \subseteq N$ not containing p_i or p_j , then $\phi_i = \phi_j$,
- **Additivity:** if ν is a sum of two value functions ν_1 and ν_2 , i.e., $\nu = \nu_1 + \nu_2$, then $\phi_i(\nu) = \phi_i(\nu_1) + \phi_i(\nu_2)$,
- **Dummy element:** if a player p_i has constant marginal contribution $\nu(\{p_i\})$ for all coalitions, i.e., $\nu(S \cup \{p_i\}) = \nu(S) + \nu(\{p_i\})$ for all $S \subseteq N \setminus \{p_i\}$, then $\phi_i = \nu(\{p_i\})$.

2.2. Multiple Arms Identification

A *multi-armed bandit* problem is specified by a set of *arms* $\mathcal{A} = \{a_1, \dots, a_n\}$ each arm a_i of which is endowed with an unknown distribution ζ_i having mean μ_i . In each discrete *time step* t , the learner can *pull* an arm a_i of its choice, meaning that it retrieves a random sample $X_i^t \sim \zeta_i$ drawn independently conditioned on the history of the previous time steps. The arms can be ordered (not necessarily uniquely) via a permutation $\pi : [n] \rightarrow [n]$ such that $\mu_{\pi(1)} \geq \dots \geq \mu_{\pi(n)}$, where we define $[n] := \{1, \dots, n\}$. Given a number $k \in [n]$, the objective of the learner in the *multiple arms identification* problem is to identify the top- k arms $a_{\pi(1)}, \dots, a_{\pi(k)}$. In the literature there are two prevalent learning frameworks for this objective, namely the *fixed budget* setting and the *fixed confidence* setting. In the former, a number of time steps T (the *budget*) is given beforehand, which once exhausted requires the learner to return its guess about the top- k arms, with its performance being measured by the probability of returning a correct output. On the contrary, the learner is judged in the latter by the number

of time steps needed in order to identify the top- k arms with probability at least $1 - \delta$ for a given $\delta \in (0, 1]$.

3. Problem Statement

The *Top- k Shapley* problem is given by a cooperative game (N, ν) in which accesses to the value function ν are costly. Although ν is known (in the sense that we can access $\nu(S)$ for all $S \subseteq N$), the Shapley values remain unknown, since it is practically infeasible for a sufficiently large number of players to compute them. The players in N can be ordered (not necessarily uniquely) via a permutation $\pi : [n] \rightarrow [n]$ such that $\phi_{\pi(1)} \geq \dots \geq \phi_{\pi(n)}$. For sake of simplicity, we assume that there are no ties at the top- k -th position. Given a number $k \in [n]$, the learner's goal is to identify the top- k players $p_{\pi(1)}, \dots, p_{\pi(k)}$ with highest Shapley values.

Likewise to multiple arms identification, we distinguish between two learning scenarios. One where performance is measured by the probability of the learner successfully identifying the top k players after a given number T of accesses to ν that the learner is allowed to make (fixed budget scenario). The other focusing on a minimal number of accesses to ν in order to guarantee a successful identification with a probability of at least $1 - \delta$ for a given $\delta \in (0, 1]$ (fixed confidence scenario). Due to page restrictions we focus only on the fixed budget setting.

4. Reduction to Multiple Arms Identification

Given a cooperative game (N, ν) , the marginal contribution $\nu(S \cup \{p_i\}) - \nu(S)$ of each player p_i can be viewed as a discrete random variable X_i if S is drawn randomly from $\mathcal{P}(N \setminus \{p_i\})$. Further, by drawing any S with probability $1/n \binom{n-1}{|S|}$, X_i has mean $\mathbb{E}[X_i] = \phi_i$. Thus, by interpreting a player p_i as an arm a_i within a multi-armed bandit problem, where retrieving a sample of the arm's distribution corresponds to drawing a (independent) sample of X_i , we obtain that the arm's mean μ_i equals the player's Shapley value ϕ_i . Together with the Shapley values, the corresponding arms' means remain unknown to us. With this connection at hand, the reduction to multiple arms identification is complete, as the objective of identifying the top- k players $p_{\pi(1)}, \dots, p_{\pi(k)}$ with highest Shapley values is equivalent to the task of finding the corresponding k arms $a_{\pi(1)}, \dots, a_{\pi(k)}$ having highest means. We denote the resulting bandit problem as *Shapley bandits*. This general reduction scheme allows leveraging any algorithm for multiple arms identification to the Top- k Shapley problem without affecting its internal mechanisms. Finally, it should be emphasized that each pull of an arm a_i involves two accesses to the value function ν , one for $\nu(S)$ and the other for $\nu(S \cup \{p_i\})$.

5. Algorithms

We present and analyze in Section 5.1 *Uniform Random Sampling* as a first benchmark algorithm, show in Section 5.2 how to adapt already existing algorithms for multiple arms identification to the top- k Shapley problem at the example of the *Gap-based Exploration* algorithm [14], and propose in Section 5.3 with *Border Uncertainty Sampling* a new algorithm that can be easily generalized to multiple arms identification.

5.1. Uniform Random Sampling

As an illustrative example of how the approach can be applied we present the *Uniform Random Sampling* algorithm (see Algorithm 1). It is a modification of the *ApproShapley* algorithm in [7] and the *Simple Random Sampling* algorithm in [8], which instead of sampling permutations of players and computing marginal contributions in the sequence in which players in the permutations appear, simply samples a coalition for each player in order to remain faithful to our reduction explained above (cf. Section 4).

For each player p_i a mean estimate $\hat{\phi}_i$ of ϕ_i is kept by URS and at termination the k players with highest estimates are returned. Note how URS does not rely on a budget T or confidence $1 - \delta$ to be given, instead it can be run for an arbitrary number of time steps and is therefore applicable for the fixed budget setting as well as the fixed confidence setting. Utilizing the

Algorithm 1: Uniform Random Sampling (URS)

Input: N, ν, k
1 Initialize: $\hat{\phi}_i \leftarrow 0, t_i \leftarrow 0 \forall p_i \in N$
2 for $t = 1, 2, \dots$ **do**
3 $i \leftarrow (t \bmod n) + 1$
4 $t_i \leftarrow t_i + 1$
5 $\phi_{i,t_i} = \nu(S \cup \{p_i\}) - \nu(S)$ with $S \subseteq N \setminus \{p_i\}$ drawn with probability $1/n \binom{n-1}{|S|}$
6 $\hat{\phi}_i \leftarrow \frac{(t_i-1)\hat{\phi}_i + \phi_{i,t_i}}{t_i}$
7 end
Output: $p_{\hat{\pi}(1)}, \dots, p_{\hat{\pi}(k)}$ for $\hat{\pi} : [n] \rightarrow [n]$ with $\hat{\phi}_{\hat{\pi}(1)} \geq \dots \geq \hat{\phi}_{\hat{\pi}(n)}$

techniques presented in [8], we can derive performance guarantees for the fixed budget and the fixed confidence setting depending on the variances or ranges of the marginal contributions of each player, stated in the following.

Theorem 1.

Let $\sigma^2 \geq \mathbb{V}[X_i]$ for all $p_i \in N$ and $k \in [n]$, $m \in \mathbb{N}$, $\delta \in (0, 1]$, as well as $\varepsilon_k > 0$ with $\varepsilon_k \leq \phi_{\pi(k)} - \phi_{\pi(k+1)}$. Then, URS identifies the top- k players correctly

- after $2mn$ many accesses to ν with probability at least $1 - 4n\sigma^2/\varepsilon_k^2 m$;
- with probability at least $1 - \delta$ after $8n^2\sigma^2/\varepsilon_k^2\delta$ many accesses to ν .

The proof is given in Appendix A. The first property becomes a guarantee for the fixed budget scenario by setting m (denoting the number of marginal contributions drawn for each player) to the highest integer fulfilling $2mn \leq T$ for the given budget T . The second property reveals a sampling complexity of $8n^2\sigma^2/\varepsilon_k^2\delta$ for the fixed confidence scenario.

Theorem 2.

Let r be an upper bound for the range of X_i for all $p_i \in N$. Further, let $k \in [n]$, $m \in \mathbb{N}$, $\delta \in (0, 1]$, and $0 < \varepsilon_k \leq \phi_{\pi(k)} - \phi_{\pi(k+1)}$. Then, URS identifies the top- k players correctly

- after $2mn$ many accesses to ν with probability at least $1 - 2n \exp(-\varepsilon_k^2 m/2r^2)$;

- with probability at least $1 - \delta$ after $4nr^2/\varepsilon_k^2 \cdot \log(2n/\delta)$ many accesses to ν .

The proof is given in Appendix B. Again, m is to be interpreted as the number of marginal contributions drawn for each player.

5.2. Gap-based Exploration

At the example of the *Gap-based Exploration* algorithm (Gap-E) [15, 14] we demonstrate how to adapt a multiple arms identification algorithm to the Top- k Shapley problem (see Algorithm 2). Originally, Gap-E was proposed and analyzed for the setting of finding the single arm with highest mean reward in [15], and later slightly modified for the task of finding the top- k arms in [14]. Whenever Gap-E pulls an arm a_i , we replace the random sample by $\nu(S \cup \{i\}) - \nu(S)$ for $S \subseteq N \setminus \{p_i\}$ drawn randomly with probability $1/n \binom{n-1}{|S|}$. Gap-E demands the budget T , a coefficient $c \in \mathbb{R}_{>0}$, and the complexity of the problem $H^{(k)}$ as additional parameters to be given, where

$$H^{(k)} = \sum_{i=1}^n \left(\Delta_i^{(k)} \right)^{-2}, \quad \text{and} \quad \Delta_i^{(k)} = \begin{cases} \mu_i - \mu_{\pi(k+1)}, & i \in \{\pi(1), \dots, \pi(k)\} \\ \mu_{\pi(k)} - \mu_i, & i \in \{\pi(k+1), \dots, \pi(n)\} \end{cases}.$$

Algorithm 2: Gap-based Exploration (Gap-E)

Input: $N, \nu, T, c, H^{(k)}$

- 1 **Initialize:** $\hat{\phi}_i \leftarrow 0, t_i \leftarrow 1 \forall p_i \in N$
- 2 **for** $i = 1, \dots, n$ **do**
- 3 $\hat{\phi}_i = \nu(S \cup \{p_i\}) - \nu(S)$ with $S \subseteq N \setminus \{p_i\}$ drawn with probability $1/n \binom{n-1}{|S|}$
- 4 **end**
- 5 **for** $t = n + 1, \dots, T$ **do**
- 6 Compute $\hat{\pi} : [n] \rightarrow [n]$ with $\hat{\phi}_{\hat{\pi}(1)} \geq \dots \geq \hat{\phi}_{\hat{\pi}(n)}$
- 7 $\Delta_i = \begin{cases} \hat{\phi}_i - \hat{\phi}_{\hat{\pi}(k+1)} & i \in \{\hat{\pi}(1), \dots, \hat{\pi}(k)\} \\ \hat{\phi}_{\hat{\pi}(k)} - \hat{\phi}_i & i \in \{\hat{\pi}(k+1), \dots, \hat{\pi}(n)\} \end{cases} \forall p_i \in N$
- 8 $i \leftarrow \arg \max_{j \in [n]} -\Delta_j + c \sqrt{\frac{T}{H^{(k)} t_j}}$
- 9 $t_i \leftarrow t_i + 1$
- 10 $\phi_{i, t_i} = \nu(S \cup \{p_i\}) - \nu(S)$ with $S \subseteq N \setminus \{p_i\}$ drawn with probability $1/n \binom{n-1}{|S|}$
- 11 $\hat{\phi}_i \leftarrow \frac{(t_i-1)\hat{\phi}_i + \phi_{i, t_i}}{t_i}$
- 12 **end**

Output: $p_{\hat{\pi}(1)}, \dots, p_{\hat{\pi}(k)}$ for $\hat{\pi} : [n] \rightarrow [n]$ with $\hat{\phi}_{\hat{\pi}(1)} \geq \dots \geq \hat{\phi}_{\hat{\pi}(n)}$

5.3. Border Uncertainty Sampling

Next, we propose a new algorithm (cf. Algorithm 3) called *Border Uncertainty Sampling* (BUS) without providing theoretical guarantees. In similar fashion to Gap-E a measure of (un-)certainty

whether a player p_i belongs to the top- k players or not is at the heart of BUS. However, the gaps involved in the measure of (un-)certainty are calculated in a slightly different manner, namely as the absolute distance to the average of the k -th and $(k + 1)$ -th highest mean estimates $\hat{\phi}_{\hat{\pi}(k)}$ and $\hat{\phi}_{\hat{\pi}(k+1)}$. Next, BUS chooses to draw a sample for the player p_i that minimizes its gap times the number of samples BUS has already drawn for it, i.e., $\Delta_i \cdot t_i$. The intuition behind this measure of certainty is that for players with larger gap Δ_i we are more certain to tell whether it belongs to the top- k players or not. Likewise, a larger number t_i of samples drawn indicates a higher precision of the estimate $\hat{\phi}_i$. Thus, BUS selects the player p_i with highest uncertainty. As with URS, a clear advantage of BUS over Gap-E is that no additional parameters like the time budget for instance are required, allowing it to be terminated at any time step.

Algorithm 3: Border Uncertainty Sampling (BUS)

Input: N, ν, k

- 1 **Initialize:** $\hat{\phi}_i \leftarrow 0, t_i \leftarrow 1 \forall p_i \in N$
- 2 **for** $i = 1, \dots, n$ **do**
- 3 $\hat{\phi}_i = \nu(S \cup \{p_i\}) - \nu(S)$ with $S \subseteq N \setminus \{p_i\}$ drawn with probability $1/n \binom{n-1}{|S|}$
- 4 **end**
- 5 **for** $t = n + 1, \dots$ **do**
- 6 Compute $\hat{\pi} : [n] \rightarrow [n]$ with $\hat{\phi}_{\hat{\pi}(1)} \geq \dots \geq \hat{\phi}_{\hat{\pi}(n)}$
- 7 $\hat{\phi}^* \leftarrow \frac{\hat{\phi}_{\hat{\pi}(k)} + \hat{\phi}_{\hat{\pi}(k+1)}}{2}$
- 8 $\Delta_i \leftarrow |\hat{\phi}_i - \hat{\phi}^*| \forall p_i \in N$
- 9 $i \leftarrow \arg \min_{j \in [n]} \Delta_j \cdot t_j$
- 10 $t_i \leftarrow t_i + 1$
- 11 $\phi_{i,t_i} = \nu(S \cup \{p_i\}) - \nu(S)$ with $S \subseteq N \setminus \{p_i\}$ drawn with probability $1/n \binom{n-1}{|S|}$
- 12 $\hat{\phi}_i \leftarrow \frac{(t_i-1)\hat{\phi}_i + \phi_{i,t_i}}{t_i}$
- 13 **end**

Output: $p_{\hat{\pi}(1)}, \dots, p_{\hat{\pi}(k)}$ for $\hat{\pi} : [n] \rightarrow [n]$ with $\hat{\phi}_{\hat{\pi}(1)} \geq \dots \geq \hat{\phi}_{\hat{\pi}(n)}$

6. Experiments

In the following we evaluate the algorithms URS, BUS, Gap-E [14], and *Successive Accepts and Rejects* (SAR) [14] modified for the Top- k Shapley problem on synthetic data. For Gap-E we have heuristically set $H^{(k)} = 10000$ and $c = 1$. We are interested in the performance curves in dependence of the number of players n , the budget T , and the variance in marginal contributions. Generating random value functions is not suitable for our purpose, as this leads to expensive computations of the corresponding Shapley values. As a remedy, we simulated cooperative games with the following two approaches. First, we consider in Section 6.1 a stochastic setting in which the marginal contributions of each player are sampled from some fixed distributions. And secondly, we simulate in 6.2 a special case of cooperative games called *sum of unanimity*

games for which the computation of Shapley values is fairly straightforward. We show in all figures for each choice of parameters the averaged ratio of correctly identified top- k players gathered from 500 repetitions.

6.1. Stochastic Setting

We substitute the marginal contributions of each player p_i by a random variable $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and set $\mu_i = 0.806 - 0.006i$ for all $p_i \in N$. The results are shown in Figure 1. For all three considered dependencies (budget, number of players, and variances) BUS outperforms the other considered algorithms by a visible margin. The performance of all algorithms improves for increasing budgets and decreasing variances as one would expect, but the impact of the number of players on BUS's and Gap-E's ratio is surprisingly low in the considered ranges.

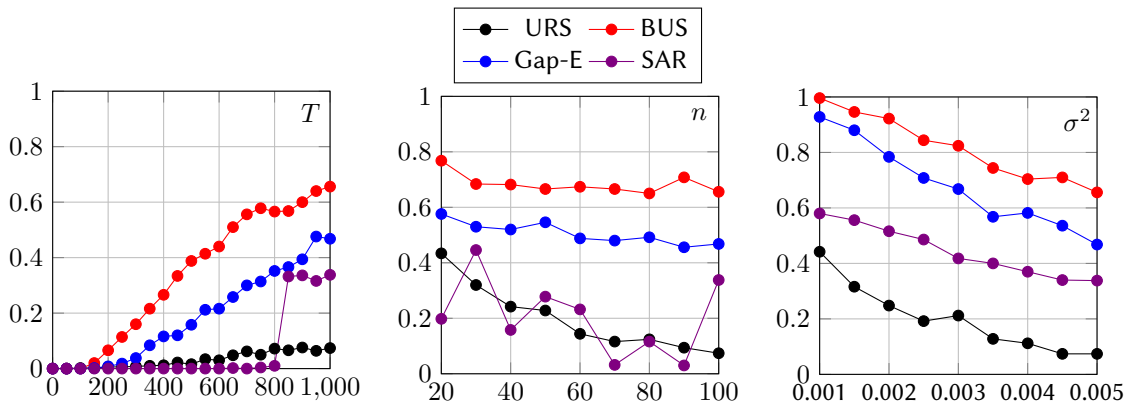


Figure 1: Averaged ratios of correct returned sets under the stochastic setting for $k = 10$. Left: $n = 100$, $\sigma^2 = 0.005$. Center: $T = 1000$, $\sigma^2 = 0.005$. Right: $n = 100$, $T = 1000$.

6.2. Sums of Unanimity Games

In an *unanimity game*, specified by a subset $R \subseteq N$, the value function takes the form of

$$\nu_R(S) = \mathbb{I}\{R \subseteq S\} \text{ for all } S \subseteq N,$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function. An unanimity game can be interpreted as a game in which all players contained in R have to agree on cooperating together in order to achieve a benefit of 1. One can construct a *sum-of-unanimity-games game* (SOUG game) by combining multiple unanimity games in a linear combination. More precisely, for a set of coalitions $\mathcal{R} \subseteq \mathcal{P}(N)$ and coefficients $c_R \in \mathbb{R}$ for each $R \in \mathcal{R}$ the value function is given by:

$$\nu(S) = \sum_{R \in \mathcal{R}} c_R \cdot \nu_R(S) \text{ for all } S \subseteq N.$$

The Shapley values of a SOUG game can be calculated in linear time with respect to the number of combined unanimity games and is given for each player p_i by [3]:

$$\phi_i = \sum_{R \in \mathcal{R}: i \in R} \frac{c_R}{|R|}.$$

For our simulations we generate SOUG games by drawing all the key terms uniformly at random within a specific range/domain, respectively. The considered ranges or domains are

- $\{5, 6, \dots, 50\}$ for the number of combined unanimity games $|\mathcal{R}|$,
- $\{0, 1, \dots, n\}$ for the size of each $R \in \mathcal{R}$,
- N for the members of each $R \in \mathcal{R}$,
- $[0, 1/|\mathcal{R}|]$ for the coefficient c_R for each $R \in \mathcal{R}$.

The results in Figure 2 show a similar picture as for the stochastic setting, albeit the performance ratios being closer together. BUS still outperforms its competitors Gap-E and SAR, while the benchmark algorithm URS does not perform significantly worse, which indicates the increased challenge that SOUG games pose in comparison to the stochastic setting. In contrast, the number of players has now a more drastic impact.




Figure 2: Averaged ratios of correct returned sets for SOUG games for $k = 3$. Left: $n = 10$. Right: $T=2000$.

7. Conclusion

We have proposed the Top- k Shapley problem, which consists of finding the k players in a cooperative game with the highest Shapley values. Taking a probabilistic view by seeing the marginal contributions of the players as discrete random variables allowed us to draw a connection to multi-armed bandits and reduce the problem to multiple-arms identification, which we have done by successfully adapting known algorithms. We proposed with BUS a new algorithm that is not limited to the use case of identifying top- k Shapley players and gave evidence for its superiority by means of empirical results. Further, it has the advantage of not

needing to know any additional parameters compared to other algorithms for multiple arms identification. For future work, we aim to derive theoretical guarantees, albeit leaving room for modifications open in order to make the analysis feasible.

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.

References

- [1] S. B. Cohen, G. Dror, E. Ruppin, Feature selection via coalitional game theory, *Neural Comput.* 19 (2007) 1939–1961.
- [2] K. Pfannschmidt, E. Hüllermeier, S. Held, R. Neiger, Evaluating tests in medical diagnosis: Combining machine learning with game-theoretical concepts, in: *IPMU*, volume 610, Springer, 2016, pp. 450–461.
- [3] L. S. Shapley, 17. A Value for n-Person Games, Princeton University Press, 2016, pp. 307–318.
- [4] T. van Campen, H. Hamers, B. Husslage, R. Lindelauf, A new approximation method for the shapley value applied to the WTC 9/11 terrorist attack, *Soc. Netw. Anal. Min.* 8 (2018) 3:1–3:12.
- [5] R. Lindelauf, H. Hamers, B. Husslage, Cooperative game theoretic centrality analysis of terrorist networks: The cases of jemaah islamiyah and al qaeda, *Eur. J. Oper. Res.* 229 (2013) 230–238.
- [6] A. Ghorbani, J. Y. Zou, Neuron shapley: Discovering the responsible neurons, in: *Advances in Neural Information Processing Systems 33*, 2020.
- [7] J. Castro, D. Gómez, J. Tejada, Polynomial calculation of the shapley value based on sampling, *Comput. Oper. Res.* 36 (2009) 1726–1730.
- [8] S. Maleki, L. Tran-Thanh, G. Hines, T. Rahwan, A. Rogers, Bounding the estimation error of sampling-based shapley value approximation with/without stratifying, *CoRR abs/1306.4265* (2013).
- [9] T. Lattimore, C. Szepesvári, *Bandit Algorithms*, Cambridge University Press, 2020.
- [10] S. Kalyanakrishnan, P. Stone, Efficient selection of multiple bandit arms: Theory and practice, in: *Proceedings of the 27th International Conference on Machine Learning*, Omnipress, 2010, pp. 511–518.
- [11] S. Kalyanakrishnan, A. Tewari, P. Auer, P. Stone, PAC subset selection in stochastic multi-armed bandits, in: *Proceedings of the 29th International Conference on Machine Learning*, icml.cc / Omnipress, 2012.
- [12] S. Chen, T. Lin, I. King, M. R. Lyu, W. Chen, Combinatorial pure exploration of multi-armed bandits, in: *Advances in Neural Information Processing Systems 27*, 2014, pp. 379–387.
- [13] Y. Zhou, X. Chen, J. Li, Optimal PAC multiple arm identification with applications to

crowdsourcing, in: Proceedings of the 31th International Conference on Machine Learning, volume 32, JMLR.org, 2014, pp. 217–225.

- [14] S. Bubeck, T. Wang, N. Viswanathan, Multiple identifications in multi-armed bandits, in: Proceedings of the 30th International Conference on Machine Learning, volume 28, JMLR.org, 2013, pp. 258–265.
- [15] V. Gabillon, M. Ghavamzadeh, A. Lazaric, S. Bubeck, Multi-bandit best arm identification, in: Advances in Neural Information Processing Systems 24, 2011, pp. 2222–2230.

A. Proof of Theorem 1

For all i and t_i we can view ϕ_{i,t_i} as a discrete random variable with:

$$\begin{aligned}\mathbb{E}[\phi_{i,t_i}] &= \sum_{S \subseteq N \setminus \{i\}} \frac{1}{n \cdot \binom{n-1}{|S|}} \cdot \nu(S \cup \{i\}) - \nu(S) \\ &= \phi_i.\end{aligned}$$

Let T_i be the number of times marginal contributions have been drawn for i and $Y_i = \sum_{t_i=1}^{T_i} \phi_{i,t_i}$, thus $\mathbb{E}[Y_i] = T_i \phi_i$ and $\hat{\phi}_i = Y_i/T_i$ at the point of termination.

Lemma 3.

Let $\varepsilon_k > 0$ with $\varepsilon_k \leq \phi_{\pi(k)} - \phi_{\pi(k+1)}$. The probability of URS identifying the top- k Shapley players correctly is at least

$$1 - \sum_{i=1}^n \mathbb{P} \left(|\hat{\phi}_{\pi(i)} - \phi_{\pi(i)}| \geq \frac{\varepsilon_k}{2} \right).$$

Proof:

First, we show that a correct identification of the top- k players by URS implies that all Shapley values are estimated with an absolute error of at most $\frac{\varepsilon_k}{2}$:

$$\begin{aligned}& \bigcup_{i=1}^k \bigcup_{j=k+1}^n \left\{ \hat{\phi}_{\pi(i)} \leq \hat{\phi}_{\pi(j)} \right\} \\ &= \bigcup_{i=1}^k \bigcup_{j=k+1}^n \left\{ \left(\hat{\phi}_{\pi(j)} - \phi_{\pi(j)} \right) + \left(\phi_{\pi(i)} - \hat{\phi}_{\pi(i)} \right) \geq \phi_{\pi(i)} - \phi_{\pi(j)} \right\} \\ &\subseteq \bigcup_{i=1}^k \bigcup_{j=k+1}^n \left\{ \hat{\phi}_{\pi(j)} - \phi_{\pi(j)} \geq \frac{\phi_{\pi(i)} - \phi_{\pi(j)}}{2} \right\} \cup \left\{ \phi_{\pi(i)} - \hat{\phi}_{\pi(i)} \geq \frac{\phi_{\pi(i)} - \phi_{\pi(j)}}{2} \right\} \\ &\subseteq \bigcup_{i=1}^k \bigcup_{j=k+1}^n \left\{ |\hat{\phi}_{\pi(j)} - \phi_{\pi(j)}| \geq \frac{\varepsilon_k}{2} \right\} \cup \left\{ |\hat{\phi}_{\pi(i)} - \phi_{\pi(i)}| \geq \frac{\varepsilon_k}{2} \right\}\end{aligned}$$

$$= \bigcup_{i=1}^n \left\{ |\hat{\phi}_{\pi(i)} - \phi_{\pi(i)}| \geq \frac{\varepsilon_k}{2} \right\}.$$

From which we derive:

$$\mathbb{P} \left(\bigcap_{i=1}^k \bigcap_{j=k+1}^n \left\{ \hat{\phi}_{\pi(i)} > \hat{\phi}_{\pi(j)} \right\} \right) \geq 1 - \sum_{i=1}^n \mathbb{P} \left(|\hat{\phi}_{\pi(i)} - \phi_{\pi(i)}| \geq \frac{\varepsilon_k}{2} \right). \quad \square$$

Let $\sigma_i^2 = \mathbb{V}[\phi_{i,t_i}]$ and hence $\mathbb{V}[Y_i] = T_i \sigma_i^2$. Similar to [8], we obtain by using Chebyshev's inequality for all $\varepsilon_k > 0$:

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \frac{\varepsilon_k}{2} \right) \leq \frac{4\sigma_i^2}{\varepsilon_k^2 T_i}.$$

We complete the proof by deriving for $\sigma \geq \sigma_i$ and $m \leq T_i$ for all i with the help of Lemma 3:

$$\mathbb{P} \left(\bigcap_{i=1}^k \bigcap_{j=k+1}^n \left\{ \hat{\phi}_{\pi(i)} > \hat{\phi}_{\pi(j)} \right\} \right) \geq 1 - \frac{4n\sigma^2}{\varepsilon_k^2 m}.$$

B. Proof of Theorem 2

Let r_i be the range of ϕ_{i,t_i} for all i . Similar to [8], we obtain by using Hoeffding's inequality, for all $\varepsilon_k > 0$:

$$\mathbb{P} \left(|\hat{\phi}_i - \phi_i| \geq \frac{\varepsilon_k}{2} \right) \leq 2 \exp \left(-\frac{\varepsilon_k^2 T_i}{2r_i^2} \right).$$

We complete the proof by deriving for $r \geq r_i$ and $m \leq T_i$ for all i with the help of Lemma 3:

$$\mathbb{P} \left(\bigcap_{i=1}^k \bigcap_{j=k+1}^n \left\{ \hat{\phi}_{\pi(i)} > \hat{\phi}_{\pi(j)} \right\} \right) \geq 1 - 2n \exp \left(-\frac{\varepsilon_k^2 M}{2r^2} \right).$$