THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The Shapley Value of Classifiers in Ensemble Games

Citation for published version:

Rozemberczki, B & Sarkar, R 2021, The Shapley Value of Classifiers in Ensemble Games. in *CIKM '21: Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. Association for Computing Machinery (ACM), pp. 1558–1567, 30th ACM International Conference on Information and Knowledge Management, Gold Coast, Queensland, Australia, 1/11/21.
<https://doi.org/10.1145/3459637.3482302>

Digital Object Identifier (DOI):

[10.1145/3459637.3482302](https://doi.org/10.1145/3459637.3482302)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

CIKM '21: Proceedings of the 30th ACM International Conference on Information & Knowledge Management

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The Shapley Value of Classifiers in Ensemble Games

Benedek Rozemberczki
The University of Edinburgh
Edinburgh, United Kingdom
benedek.rozemberczki@ed.ac.uk

Rik Sarkar
The University of Edinburgh
Edinburgh, United Kingdom
rsarkar@inf.ed.ac.uk

ABSTRACT

What is the value of an individual model in an ensemble of binary classifiers? We answer this question by introducing a class of transferable utility cooperative games called *ensemble games*. In machine learning ensembles, pre-trained models cooperate to make classification decisions. To quantify the importance of models in these ensemble games, we define *Troupe* – an efficient algorithm which allocates payoffs based on approximate Shapley values of the classifiers. We argue that the Shapley value of models in these games is an effective decision metric for choosing a high performing subset of models from the ensemble. Our analytical findings prove that our Shapley value estimation scheme is precise and scalable; its performance increases with size of the dataset and ensemble. Empirical results on real world graph classification tasks demonstrate that our algorithm produces high quality estimates of the Shapley value. We find that Shapley values can be utilized for ensemble pruning, and that adversarial models receive a low valuation. Complex classifiers are frequently found to be responsible for both correct and incorrect classification decisions.

ACM Reference Format:

Benedek Rozemberczki and Rik Sarkar. 2018. The Shapley Value of Classifiers in Ensemble Games. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

The advent of black box machine learning models raised fundamental questions about how input features and individual training data points contribute to the decisions of expert systems [17, 28]. There has also been interest in how the heterogeneity of models in an ensemble results in heterogeneous contributions of those to the classification decisions of the ensemble [16, 47]. For example one would assume that computer vision, credit scoring and fraud detection systems which were trained on varying quality proprietary datasets output labels for data points with varying accuracy. Another source of varying model performance can be the complexity of models e.g. the number of weights in a neural network or the depth of a classification tree.

Quantifying the contributions of models to an ensemble is paramount for practical reasons. Given the model valuations, the gains

of the task can be attributed to specific models, large ensembles can be reduced to smaller ones without losing accuracy [22, 32] and performance heterogeneity of ensembles can be gauged [16]. This raises the natural question: How can we measure the contributions of models to the decisions of the ensemble in an efficient, model type agnostic, axiomatic and data driven manner?

Figure 1: An overview of the model valuation problem. Models in the ensemble receive a set of data points and score those. Using the predictions and ground truth labels the oracle quantifies the worth of models.

We frame this question as one of valuation of models in an ensemble. The solution to the problem requires an analytical framework to assess the worth of individual classifiers in the ensemble. This idea is described in Figure 1. Each classifier in the ensemble receives the data points, and they output for each data point a probability distribution over the potential classes. Using these propensities an oracle – which has access to the ground truth – quantifies the worth of models in the ensemble. These importance metrics can be used to make decisions – e.g. pruning the ensemble and allocation of payoffs.

Present work. We introduce *ensemble games*, a class of transferable utility cooperative games [35]. In these games binary classifiers which form an ensemble play a voting game to assign a binary label to a data point by utilizing the features of the data point. Building on the ensemble games we derive *dual ensemble games* in which the classifiers cooperate in order to misclassify a data point. We do this to characterize the role of models in incorrect decisions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

We argue that the Shapley value [41], a solution concept from cooperative game theory, is a model importance metric. The Shapley value of a classifier in the ensemble game defined for a data point can be interpreted as the probability of the model becoming the pivotal voter in a uniformly sampled random permutation of classifiers. Computing the exact Shapley values in an ensemble game would take factorial time in the number of classifiers. In order to alleviate this we exploit an accurate approximation algorithm of the individual Shapley values which was tailored to voting games [12]. We propose *Troupe*, an algorithm which approximates the average of Shapley values in ensemble games and dual games using data. We utilize the average Shapley values as measures of model importance in the ensemble and discuss that the Shapley values are interpretable as an importance distribution over classifiers.

We evaluate *Troupe* by performing various classification tasks. Using data from real world webgraphs (Reddit, GitHub, Twitch) we demonstrate that *Troupe* outputs high quality estimates of the Shapley value. We validate that the Shapley value estimates of *Troupe* can be used as a decision metric to build space efficient and accurate ensembles. Our results establish that more complex models in an ensemble have a prime role in both correct and incorrect decisions.

Main contributions. Specifically the contributions of our work can be summarized as:

- (1) We propose ensemble games and their dual games to model the contribution of individual classifiers to the decisions of voting based ensembles.
- (2) We design *Troupe* an approximate Shapley value based algorithm to quantify the role of classifiers in decisions.
- (3) We provide a probabilistic bound for the approximation error of average Shapley values estimated from labeled data.
- (4) We empirically evaluate *Troupe* for model valuation and forward ensemble building on graph classification tasks.

The rest of this work has the following structure. In Section 2 we discuss related work on the Shapley value, its approximations and applications in machine learning. We introduce the concept of ensemble games in Section 3 and discuss Shapley value based model valuation in Section 4 with theoretical results. We evaluate the proposed algorithm experimentally in Section 5. We summarize our findings in Section 6 and discuss future work. The reference implementation of *Troupe* is available at <https://github.com/username/repositoryname> [Anonymized for double blind review].

2 RELATED WORK

Our model pruning framework intersects with research about the Shapley value and existing approaches to ensemble pruning.

2.1 The Shapley value

The *Shapley* value [41] is a solution to the problem of distributing the gains among players in a transferable utility cooperative game [35]. It is widely known for its desirable axiomatic properties [6] such as *efficiency* and *linearity*. However, exact computation of Shapley value takes factorial time, making it intractable in games with a large number of players. General [29, 37] and game specific [12] approximation techniques have been proposed. In Table 1, we

compare various approximation schemes with respect to certain desired properties.

Shapley values can be approximated using a Monte Carlo sampling of the permutations of players and also with a truncated Monte Carlo sampling variant *TMC* [29, 30, 53]. A more tractable approximation is proposed in [37], using a multilinear extension (*MLE*) of the Shapley value. A variant of this technique [23, 24] calculates the value of large players explicitly and applies the *MLE* technique to small ones. The only approximation technique tailored to weighted voting games is the expected marginal contributions method (*EMC*) which estimates the Shapley values based on contributions to varying size coalitions. Our proposed algorithm *Troupe* builds on *EMC*.

Table 1: Comparison of Shapley value computation and approximation techniques in terms of having (✓) and missing (✗) desiderata; complexities with respect to the number of players m and permutations p .

Method	Voting	Bound	Non-Random	Space	Time
Explicit	✓	✓	✓	$O(m)$	$O(m!)$
MC [30]	✗	✓	✗	$O(m)$	$O(mp)$
TMC [17]	✗	✓	✗	$O(m)$	$O(mp)$
MLE [37]	✗	✗	✓	$O(m)$	$O(m)$
MMLE [23, 24]	✗	✗	✓	$O(m)$	$O(m!)$
EMC [12]	✓	✓	✓	$O(m)$	$O(m^2)$

Shapley value has previously been used in machine learning for measuring feature importance [8, 13, 21, 26, 27, 33, 38]. In the feature selection setting the features are seen as players that cooperate to achieve high goodness of fit. Various discussed approximation schemes [9, 29, 43] have been exploited to make feature importance quantification in high dimensional spaces feasible [10, 28, 44, 45, 50] when explicit computation is not tractable. Another machine learning domain for applying the Shapley value was the pruning of neural networks [1, 18, 42]. In this context approximate Shapley values of hidden layer neurons are used to downsize overparametrized classifiers. It is argued in [42] that pruning neurons is analogous to feature selection on hidden layer features. Finally, there has been increasing interest in the equitable valuation of data points with game theoretic tools [17, 20]. In such settings the estimated Shapley values are used to gauge the influence of individual points on a supervised model. These approximate scores are obtained with group testing of features [20] and permutation sampling [5, 17].

2.2 Ensemble pruning and building

We compare *Troupe* to existing ensemble pruning and building approaches in terms of desired properties in Table 2. As one can see our framework is the only one which has all of these characteristics. Specifically, these properties are:

- *Agnostic*: The procedure can prune/build ensembles of heterogeneous model types, not just a specific type (e.g. trees).
- *Set based*: An algorithm is able to return a subset of models with a high-performance, not just a single one.
- *Diverse*: The value of individual models is implicitly or explicitly affected by how diverse their predictions are.

- *Bidirectional*: A bidirectional ensemble selection technique can select a sub-ensemble in a forward or backward manner.

Table 2: Comparison of ensemble pruning and building techniques in terms of having (✓) and missing (✗) desiderata.

Method	Agnostic	Set based	Diverse	Bidirectional
Greedy [4]	✓	✓	✗	✓
WV-LP [51]	✓	✓	✗	✗
RE [31]	✓	✓	✗	✗
DREP [25]	✓	✓	✓	✗
COMEP [2]	✗	✓	✓	✗
EP-SDP [52]	✓	✓	✗	✗
CART SySM [49]	✗	✗	✓	✗
Troupe (ours)	✓	✓	✓	✓

3 ENSEMBLES GAMES

We now introduce a novel class of co-operative games and examine axiomatic properties of solution concepts which can be applied to these games. We will discuss shapley value as an exact solution for these games, [41], and discuss approximations of the Shapley value based solution [12, 30, 37].

3.1 Ensemble game

We define an ensemble game to be one where binary classifier models (players) co-operate to label a single data point. The aggregated decision of the ensemble is assumed to be made by a vote of the players. We assume that a set of labelled data points are known:

DEFINITION 1. Labeled data point. Let $(\mathbf{x}, y)$ be a labeled data point where $\mathbf{x} \in \mathbb{R}^d$ is the feature vector and $y \in \{0, 1\}$ is the corresponding binary label.

Our work considers arbitrary binary classifier models (e.g. classification trees, support vector machines, neural networks) that operate on the same input feature vector $\mathbf{x}$. This approach is agnostic of the exact type of the model, we only assume that M can output a probability of the data point having a positive label. The model owner does not access the label, just a probability of $y = 1$ is output by the model.

DEFINITION 2. Positive classification probability for model M . Let $(\mathbf{x}, y)$ be a labeled data point and M be a binary classifier, $P(y = 1 | M, \mathbf{x})$ is the probability of the data point having a positive label output by classifier M .

We are now ready to define the operation of an ensemble classifier consisting of multiple models.

DEFINITION 3. Ensemble. An ensemble is a set $\mathcal{M}$ of size m that consists of binary classifier models $M \in \mathcal{M}$ which can each output a probability for a data point $\mathbf{x} \in \mathbb{R}^d$ having a positive label. The ensemble sets the probability of a label as:

$$P(y = 1 | \mathcal{M}, \mathbf{x}) = \sum_{M \in \mathcal{M}} P(y = 1 | M, \mathbf{x})/m.$$

And it makes decision about the label as:

$$\hat{y} = \begin{cases} 1 & \text{if } P(y = 1 | \mathcal{M}, \mathbf{x}) \geq \gamma, \\ 0 & \text{otherwise.} \end{cases}$$

where $0 \leq \gamma \leq 1$ is the decision threshold and $\hat{y}$ is the predicted label of the data point.

DEFINITION 4. Sub-ensemble. A sub-ensemble $\mathcal{S}$ is a subset $\mathcal{S} \subseteq \mathcal{M}$ of binary classifier models.

DEFINITION 5. Individual model weight. The individual weight of the vote for M , in sub-ensemble $\mathcal{S} \subseteq \mathcal{M}$ for data point $(y, \mathbf{x})$ is defined as:

$$w_M = \begin{cases} P(y = 1 | M, \mathbf{x})/m & \text{if } y = 1, \\ P(y = 0 | M, \mathbf{x})/m & \text{otherwise.} \end{cases}$$

Under this definition, the individual model weight of any binary classifier $M \in \mathcal{M}$ is bounded: $0 \leq w_M \leq 1/m$. Note that the weight of $M \in \mathcal{S}$ depends on m – the size of the larger ensemble and not on the size $|\mathcal{S}|$ of the sub-ensemble.

DEFINITION 6. Ensemble game. Let $\mathcal{M}$ be a set of binary classifiers. An ensemble game for a labeled data point $(y, \mathbf{x})$ is then a co-operative game $G = (\mathcal{M}, v)$ in which:

$$v(\mathcal{S}) = \begin{cases} 1 & \text{if } w(\mathcal{S}) \geq \gamma, \\ 0 & \text{otherwise.} \end{cases}$$

where $w(\mathcal{S}) = \sum_{M \in \mathcal{S}} w_M$ for any sub-ensemble $\mathcal{S} \subseteq \mathcal{M}$ and threshold $0 \leq \gamma \leq 1$.

This definition is the central idea in our work. The models in the ensemble play a cooperative voting game to classify the data point correctly. When the data point is classified correctly the payoff is 1, an incorrect classification results in a payoff of 0. Each model casts a weighted vote about the data point and our goal is going to be to quantify the value of individual models in the final decision. In other words, we would like to measure how individual binary classifiers contribute on average to the correct classification of a specific data point. This solution concept is described in the next section(3.2).

We can consider a misclassification as a dual ensemble game:

DEFINITION 7. Dual ensemble game. Let $\mathcal{M}$ be a set of binary classifiers. A dual ensemble game for a labeled data point $(y, \mathbf{x})$ is then a co-operative game $G = (\mathcal{M}, \tilde{v})$ in which:

$$\tilde{v}(\mathcal{S}) = \begin{cases} 1 & \text{if } \tilde{w}(\mathcal{S}) \geq \tilde{\gamma}, \\ 0 & \text{otherwise.} \end{cases}$$

for a binary classifier ensemble vote score $0 \leq \tilde{w}(\mathcal{S}) \leq 1$ where $\tilde{w}(\mathcal{S}) = \sum_{M \in \mathcal{S}} (1/m - w_M)$ for any sub-ensemble $\mathcal{S} \subseteq \mathcal{M}$ and inverse cutoff value $0 \leq \tilde{\gamma} \leq 1$ defined by $\tilde{\gamma} = 1 - \gamma$.

If the sum of classification weights for the binary classifiers is below the cutoff value the models in the ensemble misclassify the point, lose the ensemble game and as a consequence receive a payoff that is zero. In such scenarios it is interesting to ask: how can we describe the role of models in the misclassification? The dual ensemble game is derived from the original ensemble game in order to characterize this situation.

The classification game and its dual can be reframed simply as:

DEFINITION 8. Simplified ensemble game. An ensemble game in simplified form is described by the cutoff value – weight-vector tuple $(\gamma, [w_1, \dots, w_m])$.

DEFINITION 9. Simplified dual ensemble game. Given a simplified form ensemble game $(\gamma, [w_1, \dots, w_m])$, the corresponding simplified dual ensemble game is defined by the cutoff value – weight vector tuple:

$$(\tilde{\gamma}, [\tilde{w}_1, \dots, \tilde{w}_m]) = (1 - \gamma, [1/m - w_1, \dots, 1/m - w_m])$$

The simplified forms of ensemble and dual ensemble games are compact data structures which can describe the game without the models themselves and the enumeration of every sub-ensemble.

3.2 Solution concepts for model valuation

We have defined the binary classification problem with an ensemble as a weighted voting game, which is a type of co-operative game. Now we will argue that *solution concepts* of co-operative games are suitable for the valuation of individual models which form the binary classifier ensemble.

DEFINITION 10. Solution concept. A solution concept defined for the ensemble game $G = (\mathcal{M}, v)$ is a function which assigns the real value $\Phi_M(\mathcal{M}, v) \in \mathbb{R}$ to each binary classifier $M \in \mathcal{M}$.

The scalar Φ_M can be interpreted as the value of the individual binary classifier M in the ensemble $\mathcal{M}$. In the following we discuss axiomatic properties of solution concepts which are the desiderata for model valuation functions, and the implications of the axioms in the context of model valuation in binary ensemble games.

AXIOM 1. Null classifier. A solution concept has the null classifier property if $\forall S \subseteq \mathcal{M} : v(S \cup \{M\}) = v(S)$ then $\Phi_M(\mathcal{M}, v) = 0$.

Having the null classifier property means that a binary classifier which always has a zero marginal contribution in any sub-ensemble will have a zero payoff on its own. This also implies that the classifier never casts the deciding vote to correctly classify the data point when it is added to a sub-ensemble. Conversely, in the dual ensemble game the model never contributes to the misclassification of the data point.

AXIOM 2. Efficiency. A solution concept satisfies the efficiency property if $v(\mathcal{M}) = \sum_{M \in \mathcal{M}} \Phi_M(\mathcal{M}, v)$.

That is, the value (loss or gain) of an ensemble can be split precisely into the contributed value of the constituent models.

AXIOM 3. Symmetry. A solution concept has the symmetry property if $\forall S \subseteq \mathcal{M} \setminus \{M', M''\} : v(S \cup \{M'\}) = v(S \cup \{M''\})$ implies that $\Phi_{M'}(\mathcal{M}, v) = \Phi_{M''}(\mathcal{M}, v)$.

Two binary classifiers which make equal marginal contribution to all sub-ensembles have the same value in the full ensemble.

AXIOM 4. Linearity. A solution concept has the linearity property if given any two ensemble games $G = (\mathcal{M}, v)$ and $G' = (\mathcal{M}, v')$ on the same set $\mathcal{M}$, the binary classifier M satisfies $\Phi_M(\mathcal{M}, v + v') = \Phi_M(\mathcal{M}, v) + \Phi_M(\mathcal{M}, v')$.

That is, the value in the combined game is the sum of the values in individual games. This property will imply that valuations of a model for different datapoints, when added, leads to its valuation on the dataset.

3.3 The Shapley value

The Shapley value [41] of a classifier is the average marginal contribution of the model over the possible different permutations in which the ensemble can be formed [6]. It is a solution concept which satisfies Axioms 1-4 and the only solution concept which is uniquely characterized by Axioms 3 and 4.

DEFINITION 11. Shapley value. The Shapley value of binary classifier M in the ensemble $\mathcal{M}$, for the data point level ensemble game $G = (\mathcal{M}, v)$ is defined as

$$\Phi_M(v) = \sum_{S \subseteq \mathcal{M} \setminus \{M\}} \frac{|S|! (|\mathcal{M}| - |S| - 1)!}{|\mathcal{M}|!} (v(S \cup \{M\}) - v(S)).$$

Calculating the exact Shapley value for every model in an ensemble game would take $O(m!)$ time, or more, which is computationally unfeasible in large ensembles. We discuss a range of approximation approaches in detail which can give Shapley value estimates in $O(m)$ and $O(m^2)$ time.

3.3.1 Multilinear extension (MLE) approximation of the Shapley value. The MLE approximation of the Shapley value in a voting game [12, 37] can be used to estimate the Shapley value in the ensemble game and its dual game. Let us define the expectation and the variation of the aggregated ensemble contributions for the remaining models as: $\mu_M = \sum_{j=1}^m w_j - w_M$ and $v_M = \sum_{j=1}^m w_j^2 - w_M^2$. For a classifier $M \in \mathcal{M}$ the multi-linear approximation of the unnormalized Shapley value is computed by:

$$\hat{\Phi}_M \propto \int_{-\infty}^y \frac{1}{\sqrt{2\pi v_M}} \exp\left(-\frac{(x - \mu_M)^2}{2v_M}\right) dx - \int_{-\infty}^{y - w_M} \frac{1}{\sqrt{2\pi v_M}} \exp\left(-\frac{(x - \mu_M)^2}{2v_M}\right) dx.$$

This approximation assumes that the size of the game is large (many classifiers in the ensemble in our case) and also that μ has an *approximate normal distribution*. Calculating all of the approximate Shapley values by MLE takes $O(m)$ time.

3.3.2 Monte Carlo (MC) approximation of the Shapley value. The MC approximation [29, 30] given the ensemble $\mathcal{M}$ estimates the Shapley value of the model $M \in \mathcal{M}$ by the average marginal contribution over uniformly sampled permutations.

$$\hat{\Phi}_M = \mathbb{E}_{\theta \sim \Theta} [v(S_\theta^M \cup \{M\}) - v(S_\theta^M)] \quad (1)$$

In Equation (1), Θ is a uniform distribution over the $m!$ permutations of the binary classifiers and S_θ^M is the subset of models that appear before the classifier M in permutation θ . Approximating the Shapley value requires the generation of p classifier permutations (for a suitable p), and marginal contribution calculations with respect to those contributions – this takes $O(mp)$ time.

3.3.3 Voting game approximation of the Shapley value. The ensemble games introduced above are a variant of voting games [36], hence we can use the *Expected Marginal Contributions (EMC)* approximation [12]. This procedure sums the expected marginal contributions of a model to fixed size ensembles – it is described in the pseudo-code Algorithm 1.

The algorithm iterates over the individual model weights and initializes Shapley values as zeros (lines 1-2). For each ensemble size it calculates the expected contribution of the model to the ensemble.

Figure 2: The approximate average Shapley value calculation pipeline. (a) Given the set of labeled datapoints the data owner sends the features of data points to the binary classifiers in the ensemble. (b) The models in the ensemble score each of the data points independently and send the probability scores back to the oracle. (c) Given the label and the scores the oracle defines the ensemble game on each of the data points. (d-e) The ensemble games and the dual ensemble games are solved by the oracle and approximate Shapley values of models are calculated in each of the data point level ensemble games.

This expected contribution is the probability that a classifier becomes the marginal voter. These marginal contributions are added to the approximate Shapley value (lines 3-8). Using this Shapley value approximation technique takes $O(m^2)$ time. However, it is the most accurate Shapley value approximation technique in terms of absolute error of the Shapley value estimates [12].

Data: $[w_1, \dots, w_m]$ – Weights of binary classifiers.
 γ – Cutoff value.
 δ – Numerical stability parameter.
 μ – Expected value of weights.
 v – Variance of weights.

Result: $(\hat{\Phi}_1, \dots, \hat{\Phi}_m)$ – Approximate Shapley values.

```

1 for  $j \in \{1, \dots, m\}$  do
2    $\hat{\Phi}_j \leftarrow 0$ 
3   for  $k \in \{1, \dots, m-1\}$  do
4      $a \leftarrow (\gamma - w_j/k)$ 
5      $b \leftarrow (\gamma - \delta)/k$ 
6      $\hat{\Phi}_j \leftarrow \hat{\Phi}_j + \frac{1}{\sqrt{2\pi v/k}} \int_a^b \exp(-k \frac{(x-\mu)^2}{2v}) dx$ 
7   end
8 end

```

Algorithm 1: Expected marginal contribution approximation of Shapley values based on [12].

Data: $(\mathbf{x}, y)$ – Labeled data point.
 $\{M_1, \dots, M_m\}$ – Set of binary classifiers.

Result: $[w_1, \dots, w_m]$ – Weights of individual models.

```

1 for  $j \in \{1, \dots, m\}$  do
2   if  $y = 1$  then
3      $w_j \leftarrow P(y = 1 | M_j, \mathbf{x})/m$ 
4   else
5      $w_j \leftarrow P(y = 0 | M_j, \mathbf{x})/m$ 
6   end
7 end

```

Algorithm 2: Calculating the individual model weights in an ensemble game given a data point and a classifier ensemble.

4 MODEL VALUATION WITH THE SHAPLEY VALUE OF ENSEMBLE GAMES

In this section we present *Troupe*, the mechanism used for model valuation in ensembles. The chart in Figure 2 gives a high-level summary of our model valuation approach.

4.1 The model valuation algorithm

The Shapley value based model evaluation procedure described by Algorithm 3 uses $\mathcal{D} = \{(\mathbf{x}_1, y_1); \dots; (\mathbf{x}_n, y_n)\}$, a set of n datapoints with binary labels, for evaluating the ensemble $\mathcal{M}$ given a cutoff value $0 \leq \gamma \leq 1$ and a numerical stability parameter $\delta \in \mathbb{R}^{0+}$.

We set m the number of models and n the number of labeled data points used in the model valuation (lines 1-2). We iterate over all of the data points (line 3) and given a data point for all of the models we compute the individual model weights in an ensemble game using the label, features and the model itself (line 4) – see Algorithm 2 for the exact details. Using the individual model weights in the ensemble game we calculate the expected value and variance of the data point (lines 5-6).

If the sum of individual model weights is higher than the cutoff value (line 7) the Shapley values in the ensemble game are approximated by Algorithm 1 for the data point (line 8). The Shapley values in the dual ensemble game are defined by a vector of zeros (line 9). If the point is misclassified (line 10) a *dual ensemble game* is defined for the data point. This requires the redefined cutoff value $\tilde{\gamma}$, expected value $\tilde{\mu}$ and variance $\tilde{v}$ of the individual model weight vector (lines 11-13). We exploit the linearity of expectation and the location invariance of the variance. The individual model weight vector is redefined (line 14) in order to quantify the role of models in the erroneous decision. The original ensemble game Shapley values are initialized with zeros (line 15) and using the new parametrization of the game we approximate the dual ensemble game Shapley values (line 16) by Algorithm 1. The algorithm outputs data point level Shapley values for each model.

Data: $\mathcal{D}$ – Labeled data points.
 $\mathcal{M}$ – Set of binary classifiers.
 γ – Cutoff value.
 δ – Numerical stability parameter.

Result: $\{(\widehat{\Phi}_0^{1,+}, \dots, \widehat{\Phi}_m^{1,+}), \dots, (\widehat{\Phi}_0^{n,+}, \dots, \widehat{\Phi}_m^{n,+})\}$
 Shapley value vectors for ensemble games.
 $\{(\widehat{\Phi}_0^{1,-}, \dots, \widehat{\Phi}_m^{1,-}), \dots, (\widehat{\Phi}_0^{n,-}, \dots, \widehat{\Phi}_m^{n,-})\}$
 Shapley value vectors for dual ensemble games.

```

1  $m \leftarrow |\mathcal{M}|$ 
2  $n \leftarrow |\mathcal{D}|$ 
3 for  $i \in \{1, \dots, n\}$  do
4    $[w_1^i, \dots, w_m^i] \leftarrow \text{GetWeights}((\mathbf{x}_i, y_i); \{M_1, \dots, M_m\})$ 
5    $\mu \leftarrow \sum_{j=1}^m w_j^i / m$ 
6    $v \leftarrow \sum_{j=1}^m (w_j^i - \mu)^2 / m$ 
7   if  $\mu > \gamma / m$  then
8      $(\widehat{\Phi}_0^{i,+}, \dots, \widehat{\Phi}_m^{i,+}) \leftarrow \text{Shapley}([w_1^i, \dots, w_m^i]; \gamma; \delta; \mu; v)$ 
9      $(\widehat{\Phi}_0^{i,-}, \dots, \widehat{\Phi}_m^{i,-}) \leftarrow \mathbf{0}$ 
10  else
11     $\widetilde{\gamma} \leftarrow 1 - \gamma$ 
12     $\widetilde{\mu} \leftarrow 1/m - \mu$ 
13     $\widetilde{v} \leftarrow v$ 
14     $[\widetilde{w}_1^i, \dots, \widetilde{w}_m^i] \leftarrow [1/m - w_1^i, \dots, 1/m - w_m^i]$ 
15     $(\widehat{\Phi}_0^{i,+}, \dots, \widehat{\Phi}_m^{i,+}) \leftarrow \mathbf{0}$ 
16     $(\widehat{\Phi}_0^{i,-}, \dots, \widehat{\Phi}_m^{i,-}) \leftarrow \text{Shapley}([\widetilde{w}_1^i, \dots, \widetilde{w}_m^i]; \widetilde{\gamma}; \delta; \widetilde{\mu}; \widetilde{v})$ 
17  end
18 end

```

Algorithm 3: *Troupe*: Calculating the approximate Shapley value of the classifiers in ensemble and dual ensemble games.

4.2 Measuring ensemble heterogeneity

The data point level Shapley values of ensemble games and their dual can be aggregated to measure the importance of models in ensemble level decisions.

In order to quantify the role of models in classification and misclassification we calculate the average Shapley value of models in the ensemble and dual ensemble games conditional on the success of classification. The sets $\mathcal{N}^+$ and $\mathcal{N}^-$ contain the indices of classified and misclassified data points. Using the cardinality of these sets $n^+ = |\mathcal{N}^+|$, $n^- = |\mathcal{N}^-|$ and the data point level Shapley values output by Algorithm 1 we can estimate the *conditional* role of models in classification and misclassification by averaging the approximate Shapley values using Equations (2) and (3).

$$(\overline{\Phi}_1^+, \dots, \overline{\Phi}_m^+) = \sum_{i \in \mathcal{N}^+} (\widehat{\Phi}_0^{i,+}, \dots, \widehat{\Phi}_m^{i,+}) / n^+ \quad (2)$$

$$(\overline{\Phi}_1^-, \dots, \overline{\Phi}_m^-) = \sum_{i \in \mathcal{N}^-} (\widehat{\Phi}_0^{i,-}, \dots, \widehat{\Phi}_m^{i,-}) / n^- \quad (3)$$

If a component of the average Shapley value vector in Equation (2) is large compared to other components the corresponding model has an important role in the correct classification decisions of the

ensemble. A large component in Equation (3) corresponds to a model which is responsible for a large number of misclassifications.

4.3 Theoretical properties

Our framework utilizes labeled data instances to approximate the importance of classifiers in the ensemble and this has important implications. In the following we discuss how the size of the dataset and the number classifiers affects the Shapley value approximation error and the runtime of *Troupe*.

4.3.1 Bounding the average approximation error. Our discussion focuses on the average conditional Shapley value in ensemble games. However, analogous results can be obtained for the Shapley values computed from dual ensemble games.

DEFINITION 12. Approximation error. The Shapley value approximation error of model $M \in \mathcal{M}$ in an ensemble game is defined as $\Delta\Phi_M^+ = \widehat{\Phi}_M^+ - \Phi_M^+$.

DEFINITION 13. Average approximation error. Let us denote the Shapley value approximation errors of model $M \in \mathcal{M}$ calculated from the dataset $\mathcal{D}$ of correctly classified points as $\Delta\Phi_M^{1,+}, \dots, \Delta\Phi_M^{n,+}$. The average approximation error is defined by $\overline{\Delta\Phi}_M^+ = \sum_{i=1}^n \Delta\Phi_M^{i,+} / n$.

THEOREM 1. Average conditional Shapley value error bound. If the Shapley value approximation errors $\Delta\Phi_M^{1,+}, \dots, \Delta\Phi_M^{n,+}$ of model $M \in \mathcal{M}$ calculated by Algorithm 3 from the dataset $\mathcal{D}$ are independent random variables than for any $\varepsilon \in \mathbb{R}^+$ Inequality (4) holds.

$$P(|\overline{\Delta\Phi}_M^+ - \mathbb{E}[\overline{\Delta\Phi}_M^+]| \geq \varepsilon) \leq 2 \exp \left(-\sqrt{\frac{n^2 \cdot m \cdot \varepsilon^4 \cdot \pi}{8}} \right) \quad (4)$$

PROOF. Let us first note the fact that every absolute Shapley approximation value is bounded by the inequality described in Lemma 1.

LEMMA 1. Approximate Shapley value bound. As Theorem (10) of [12] states the approximation error of the Shapley value in a single voting game is bounded by Inequality (5) when the expected marginal contributions approximation is used.

$$-\sqrt{\frac{8}{m\pi}} \leq \Delta\Phi_M^+ \leq \sqrt{\frac{8}{m\pi}} \quad (5)$$

Using Lemma 1, the fact that *Troupe* is based on the expected marginal contributions approximation and that the Shapley values of a model in different ensemble games are independent random variables we can use Hoeffding's second inequality [19] for bounded non zero-mean random variables:

$$P(|\overline{\Delta\Phi}_M^+ - \mathbb{E}[\overline{\Delta\Phi}_M^+]| \geq \varepsilon) \leq 2 \exp \left(-\frac{2\varepsilon^2 n^2}{\sum_{i=1}^n \left(\left(\sqrt{\frac{8}{m\pi}} \right) - \left(-\sqrt{\frac{8}{m\pi}} \right) \right)^2} \right)$$

$$P(|\overline{\Delta\Phi}_M^+ - \mathbb{E}[\overline{\Delta\Phi}_M^+]| \geq \varepsilon) \leq 2 \exp \left(-\frac{2\varepsilon^2 n^2}{2n \sqrt{\frac{8}{m\pi}}} \right)$$

□

THEOREM 2. Confidence interval of the expected average approximation error. In order to acquire an $(1-\alpha)$ -confidence interval of $\mathbb{E}[\Delta\Phi_M^+] \pm \epsilon$ one needs a labeled dataset $\mathcal{D}$ of correctly classified data points for which n the cardinality of $\mathcal{D}$, satisfies Inequality (6).

$$n \geq \sqrt{\frac{8 \ln^2\left(\frac{\alpha}{2}\right)}{\epsilon^4 m \pi}} \quad (6)$$

PROOF. The probability $P(|\Delta\Phi_M^+ - \mathbb{E}[\Delta\Phi_M^+]| \geq \epsilon)$ in Theorem 1 equals to the level of significance for the confidence interval $\mathbb{E}[\Delta\Phi_M^+] \pm \epsilon$. Which means that Inequality (7) holds for the significance level α .

$$\alpha \leq 2 \exp\left(-\sqrt{\frac{n^2 \cdot m \cdot \epsilon^4 \cdot \pi}{8}}\right). \quad (7)$$

Solving inequality (7) for n yields the cardinality of the dataset (number of correctly classified data points) required for obtaining the confidence interval described in Theorem 2. $\square$

The inequality presented in Theorem 2 has two important consequences regarding the bound:

- (1) Larger ensembles require less data in order to give confident estimates of the Shapley value for individual models.
- (2) The dataset size requirement is sublinear in terms of confidence level and quadratic in the precision of the Shapley value approximation.

4.3.2 Runtime and memory complexity. The runtime and memory complexity of the proposed model valuation framework depends on the complexity of the main evaluation phases. We assume that our framework operates in a single-core non distributed setting.

Scoring and game definition. Assuming that the scoring of a data point takes $O(1)$ time, scoring the data point with all models takes $O(m)$. Scoring the whole dataset and defining games both takes $O(nm)$ time and $O(nm)$ space respectively.

Approximation and overall complexity. Calculating the expected marginal contribution of a model to a fixed size ensemble takes $O(1)$ time. Doing this for all of the ensemble sizes takes $O(m)$ time. Approximating the Shapley value for all models requires $O(m^2)$ time and $O(m)$ space. Given a dataset of n points this implies a need for $O(nm^2)$ time and $O(nm)$ space. This is also the overall time and space complexity of the proposed framework.

5 EXPERIMENTAL EVALUATION

In this section, we show that *Troupe* approximates the average of Shapley values precisely. We provide evidence that Shapley values are a useful decision metric for ensemble creation. Our results illustrate that model importance and complexity are correlated, and that Shapley values are able to identify adversarial models in the ensemble. Our evaluation is based on various real world binary graph classification tasks [39]. Specifically, we use datasets collected from *Reddit*, *Twitch* and *GitHub* – the descriptive statistics of these datasets are in Table 3.

Table 3: Descriptive statistics of the graph classification datasets taken from [39] used for the evaluation of *Troupe*.

Dataset	Classes		Nodes		Density		Diameter	
	Positive	Negative	Min	Max	Min	Max	Min	Max
Reddit	521	479	11	93	0.023	0.027	2	18
Twitch	520	480	14	52	0.039	0.714	2	2
GitHub	552	448	10	942	0.004	0.509	2	15

5.1 The precision of approximation

The Shapley value approximation performance of *Troupe* is compared to that of various other estimation schemes [29, 37] discussed earlier. We use an ensemble of logistic regressions where each classifier is trained on features extracted with a whole graph embedding technique [40, 46, 48]. We utilize 50% of the graphs for training and calculated the average conditional Shapley values of ensemble games and dual ensemble games using the remaining 50% of the data.

5.1.1 Experimental details. The features of the graphs are extracted with whole graph embedding techniques implemented in the open source *Karate Club* framework [39]. Given a set of graphs $\mathcal{G} = (G_1, \dots, G_n)$ whole graph embedding algorithms [40, 46] learn a mapping $g : \mathcal{G} \rightarrow \mathbb{R}^d$ which delineate the graphs $G \in \mathcal{G}$ to a d dimensional metric space. We utilize the following whole graph embedding and statistical fingerprinting techniques:

- (1) **FEATHER** [40] uses the characteristic function of topological features as a graph level statistical descriptor.
- (2) **Graph2Vec** [34] extracts tree features from the graph.
- (3) **GL2Vec** [7] distills tree features from the dual graph.
- (4) **NetLSD** [46] derives characteristics of graphs using the heat trace of the graph spectra.
- (5) **SF** [11] utilizes the largest eigenvalues of the graph Laplacian matrix as an embedding.
- (6) **LDP** [3] sketches the histogram of local degree distributions.
- (7) **GeoScattering** [15] applies the scattering transform to various structural features (e.g. degree centrality).
- (8) **IGE** [14] combines graph features from local degree distributions and scattering transforms.
- (9) **FGSD** [48] sketches the Moore-Penrose spectrum of the normalized graph Laplacian with a histogram.

The embedding techniques use the *default settings* of the *Karate Club* library, each embedding dimension is column normalized and the graph features are fed to the *scikit-learn* implementation of logistic regression. This classifier is trained with ℓ_2 penalty cost, we choose an SGD optimizer and the regularization coefficient λ was set to be 10^{-2} .

5.1.2 Experimental findings. In Table 4 we summarize the absolute percentage error values (compared to exact Shapley values in ensemble games) obtained with the various approximations schemes. (i) Our empirical results support that *Troupe* consistently computes accurate estimates of the ground truth model evaluations across datasets and classifiers in the ensemble. (ii) The high quality of estimates suggests that the approximate Shapley values of models extracted by *Troupe* can serve as a proxy decision metric for ensemble building and model selection.

Table 4: Absolute percentage error of average conditional Shapley values obtained by approximation techniques (rows) for the graph classifiers (columns) in the ensemble game. Bold numbers note the lowest error on each dataset – classifier pair.

	Approximation	FEATHER	Graph2Vec	GL2Vec	NetLSD	SF	LDP	GeoScatter	IGE	FGSD
Reddit	Troupe	1.23	2.35	8.18	0.99	2.64	2.31	1.64	4.85	1.49
	MLE	3.20	23.61	32.62	4.19	5.34	7.97	7.12	5.42	7.61
	MC $p = 10^3$	12.57	30.94	13.26	8.67	32.76	12.32	11.62	16.36	12.78
	MC $p = 10^3$	4.71	5.38	3.41	1.67	3.03	5.51	4.34	4.97	3.82
Twitch	Troupe	0.28	3.33	1.18	2.53	1.62	0.59	1.48	0.25	1.19
	MLE	5.22	5.44	3.05	8.32	2.38	1.92	3.14	4.85	5.77
	MC $p = 10^2$	2.37	10.40	6.76	7.07	15.79	6.36	13.99	23.96	0.39
	MC $p = 10^3$	2.32	4.60	2.96	2.67	2.53	2.73	6.31	0.27	3.89
GitHub	Troupe	2.68	0.18	2.61	1.41	1.49	1.23	1.88	2.36	1.04
	MLE	9.22	5.12	3.76	3.27	9.71	7.46	5.08	4.04	0.73
	MC $p = 10^2$	5.91	9.37	4.67	9.82	8.78	8.34	13.66	28.95	0.76
	MC $p = 10^3$	3.35	6.09	7.70	3.26	2.84	0.79	6.67	2.51	1.12

5.2 Ensemble building

Our earlier results suggested that the approximate Shapley values output by *Troupe* (and other estimation methods) can be used as decision metrics for ensemble building.

Figure 3: The classification performance of ensembles as a function of ensemble size selected by our forward model building procedure and baselines.

5.2.1 Experimental settings. We demonstrate this by selecting a high performance subset of a random forest in a forward fashion using the estimated model valuation scores. From each graph we extract Weisfeiler-Lehman tree features [34] and keep those topological patterns which appeared in at least 5 graphs. Utilizing the counts of these features in graphs we define statistical descriptors. The selection procedure and evaluation steps are the following:

- (1) Using 40% of the graphs we train a random forest with 50 classification trees, each tree is trained on 20 randomly sampled

Weisfeiler-Lehman features – we use the default settings of *scikit-learn*.

- (2) We calculate the average conditional Shapley value of classifiers in the games using 30% of the data.
- (3) We order the classifiers by the approximate Shapley values in decreasing order, create subensembles in a forward fashion and calculate the predictive performance of the resulting subensembles on the remaining 30% of the graphs.

The test performance (measured by AUC scores) of these classifiers and baselines as a function of ensemble size is plotted on Figure 3.

5.2.2 Experimental findings. Our results suggest that *Troupe* has a material advantage over competing Shapley value approximation schemes. Moreover, the proposed method outperforms the selection of ensemble building baselines on the Twitch and Reddit datasets [25, 31, 51, 52]. This implies that *Troupe* is a good alternative for existing ensemble pruning techniques and our choice of the Shapley value approximation scheme is justified in practical applications.

Figure 4: The distribution of normalized Shapley values for neural networks in ensemble and dual ensemble games conditional on the number of neurons.

5.3 Model complexity and influence

The Shapley value gives an implicit measure of model influence on the ensemble. It is interesting to see whether there is a connection between relative importance and model complexity (e.g. depth of trees or number of neurons). In order to investigate this, we create a voting expert which consists of neural networks with heterogeneous model complexity.

5.3.1 Experimental settings. We create an ensemble of $m = 10^3$ neural networks using *scikit-learn* – each of these has a single hidden layer. Each model receives 20 randomly selected frequency features as input and has a randomly chosen number of hidden layer neurons – we uniformly sample this hyperparameter from $\{2^3, 2^4, 2^5, 2^6, 2^7\}$. Individual neural networks are trained by minimizing the binary cross-entropy with SGD for 200 epochs with a learning rate of 10^{-2} .

5.3.2 Experimental findings. The distribution of normalized average Shapley values for the Reddit dataset are plotted on Figure 4 for the ensemble and dual ensemble games conditioned on the number of hidden layer neurons. The results imply that more complex models with a larger number of free parameters receive higher Shapley values in both classes of games. In simple terms complex models contribute to correct and incorrect classification decisions at a disproportionate rate.

5.4 Identifying adversarial models

Ensembles can be formed by the model owners submitting their classifiers voluntarily to a machine learning market. We investigated how adversarial behaviour of model owners affects the Shapley values of classifiers.

Figure 5: The mean Shapley value (standard errors added) of adversarial and base classifiers in ensemble games as a function of adversarial noise ratio mixed to predictions.

5.4.1 Experimental settings. We use the Weisfeiler-Lehman tree features described in Subsections 5.2 and 5.3 to train a random forest ensemble of 20 classifiers using 50% of the data – each of these models utilizes a random subset of 20 features and had a maximal tree depth of 4. The Shapley values are calculated from the remaining 50% of the data using *Troupe*. We artificially corrupt the predictions for 10 of the classification trees by mixing the outputted probability values with noise that has $\mathcal{U}(0, 1)$ distribution. The corrupt predictions are a convex combination of the original prediction and the noise – we call the weight of the noise as the adversarial noise ratio. Given an adversarial noise ratio we calculate the mean Shapley value (with standard errors) for the adversarial and normally behaving classification trees in the ensemble.

5.4.2 Experimental findings. In Figure 5 we plotted the mean Shapley value of models which are adversarial and which are not in a scenario where half of the model owners mixed their predictions with noise. Even with a negligible amount of adversarial noise the Shapley values of adversarial models drop in the ensemble

games considerably. This implies that our method can be used to find adversarial models in a community contributed classification ensemble.

5.5 Scalability

We plotted the mean runtime of Shapley value approximations calculated from 10 experimental runs for an ensemble with $m = 2^5$ and dataset with $n = 2^8$ on Figure 6. All results are obtained with our open-source Shapley value approximation framework. These average runtimes of the approximation techniques are in line with the known and new theoretical results discussed in Sections 2 and 4. The precise estimates of Shapley values obtained with *Troupe* come at a runtime cost.

Figure 6: The runtime of Shapley value approximation in ensemble games as a function of dataset size and number of classifiers in the ensemble.

6 CONCLUSIONS AND FUTURE DIRECTIONS

We proposed a new class of cooperative games called *ensemble games* in which binary classifiers cooperate in order to classify a data point correctly. We postulated that solving these games with the Shapley value results in a measure of individual classifier quality. We designed *Troupe* a voting game inspired approximation algorithm which computes the average of Shapley values for every classifier in the ensemble based on a dataset. We provided theoretical results about the sample size needed for precise estimates of the model quality.

We have demonstrated that our algorithm can provide accurate estimates of the Shapley value on real world social network data. We illustrated how the Shapley values of the models can be used to create small sized but highly accurate graph classification ensembles. We presented evidence that complex models have an important role in the classification decisions of ensembles. We showcased that our framework can identify adversarial models in the classification ensemble.

We think that our contribution opens up venues for novel theoretical and applied data mining research about ensembles. We theorize that our model valuation framework can be generalized to ensembles which consist of multi-class predictors using a one versus many approach. Our work provides an opportunity for the design and implementation of real world machine learning markets where the payoff of a model owner is a function of the individual model value in the ensemble.

REFERENCES

- [1] Marco Ancona, Cengiz Oztireli, and Markus Gross. 2019. Explaining Deep Neural Networks with a Polynomial Time Algorithm for Shapley Value Approximation. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Vol. 97. PMLR, 272–281.
- [2] Yijun Bian, Yijun Wang, Yaqiang Yao, and Huanhuan Chen. 2019. Ensemble Pruning Based on Objection Maximization With a General Distributed Framework. *IEEE Transactions on Neural Networks and Learning Systems* 31, 9 (2019), 3766–3774.
- [3] Chen Cai and Yusu Wang. 2018. A Simple Yet Effective Baseline for Non-Attributed Graph Classification. *ICLR 2019 Workshop on Representation Learning on Graphs and Manifolds* (2018).
- [4] Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. 2004. Ensemble Selection from Libraries of Models. In *Proceedings of the Twenty-First International Conference on Machine Learning*. 18.
- [5] Javier Castro, Daniel Gómez, and Juan Tejada. 2009. Polynomial Calculation of the Shapley Value Based on Sampling. *Computers and Operations Research* 36, 5 (2009), 1726 – 1730.
- [6] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. Computational Aspects of Cooperative Game Theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5, 6 (2011), 1–168.
- [7] Hong Chen and Hisashi Koga. 2019. GL2Vec: Graph Embedding Enriched by Line Graphs with Edge Features. In *International Conference on Neural Information Processing*. Springer, 3–14.
- [8] Jianbo Chen, Le Song, Martin J Wainwright, and Michael I Jordan. 2018. L-Shapley and C-Shapley: Efficient Model Interpretation for Structured Data. In *International Conference on Learning Representations*.
- [9] Shay Cohen, Eytan Ruppin, and Gideon Dror. 2005. Feature Selection Based on the Shapley Value. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*. 665–670.
- [10] Ian Covert and Su-In Lee. 2021. Improving KernelSHAP: Practical Shapley Value Estimation Using Linear Regression. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3457–3465.
- [11] Nathan de Lara and Pineau Edouard. 2018. A Simple Baseline Algorithm for Graph Classification. In *Advances in Neural Information Processing Systems*.
- [12] Shaheen S Fatima, Michael Wooldridge, and Nicholas R Jennings. 2008. A Linear Approximation Method for the Shapley Value. *Artificial Intelligence* 172, 14 (2008), 1673–1699.
- [13] Christopher Frye, Colin Rowat, and Ilya Feige. 2020. Asymmetric Shapley values: incorporating causal knowledge into model-agnostic explainability. *Advances in Neural Information Processing Systems* 33 (2020).
- [14] Alexis Galland and Marc Lelarge. 2019. Invariant Embedding for Graph Classification. In *ICML Workshop on Learning and Reasoning with Graph-Structured Data*.
- [15] Feng Gao, Guy Wolf, and Matthew Hirn. 2019. Geometric Scattering for Graph Data Analysis. In *Proceedings of the 36th International Conference on Machine Learning*, Vol. 97. 2122–2131.
- [16] Mike Gashler, Christophe Giraud-Carrier, and Tony Martinez. 2008. Decision Tree Ensemble: Small Heterogeneous is Better than Large Homogeneous. In *Seventh International Conference on Machine Learning and Applications*. IEEE, 900–905.
- [17] Amirata Ghorbani and James Zou. 2019. Data Shapley: Equitable Valuation of Data for Machine Learning. In *International Conference on Machine Learning*. 2242–2251.
- [18] Amirata Ghorbani and James Zou. 2020. Neuron Shapley: Discovering the Responsible Neurons. *arXiv preprint arXiv:2002.09815* (2020).
- [19] Wassily Hoeffding. 1994. Probability Inequalities for Sums of Bounded Random Variables. In *The Collected Works of Wassily Hoeffding*. Springer, 409–426.
- [20] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J. Spanos. [n.d.]. Towards Efficient Data Valuation Based on the Shapley Value. In *Proceedings of Machine Learning Research*. 1167–1176.
- [21] Christophe Labreuche and Simon Fossier. 2018. Explaining Multi-Criteria Decision Aiding Models with an Extended Shapley Value. In *IJCAI*. 331–339.
- [22] Aleksandar Lazarevic and Zoran Obradovic. 2001. Effective Pruning of Neural Network Classifier Ensembles. In *IJCNN'01. International Joint Conference on Neural Networks. Proceedings (Cat. No. 01CH37222)*, Vol. 2. IEEE, 796–801.
- [23] Dennis Leech. 1998. *Computing Power Indices for Large Voting Games: A New Algorithm*. Technical Report.
- [24] Dennis Leech. 2003. Computing Power Indices for Large Voting Games. *Management Science* 49, 6 (2003), 831–837.
- [25] Nan Li, Yang Yu, and Zhi-Hua Zhou. 2012. Diversity Regularized Ensemble Pruning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 330–345.
- [26] Stan Lipovetsky and Michael Conklin. 2001. Analysis of Regression in Game Theory Approach. *Applied Stochastic Models in Business and Industry* 17, 4 (2001), 319–330.
- [27] Yifei Liu, Chao Chen, Yazheng Liu, Xi Zhang, and Sihong Xie. 2020. Shapley Values and Meta-Explanations for Probabilistic Graphical Model Inference. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 945–954.
- [28] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Curran Associates Inc., 4768–4777.
- [29] Sasan Maleki, Long Tran-Thanh, Greg Hines, Talal Rahwan, and Alex Rogers. 2013. Bounding the Estimation Error of Sampling-Based Shapley Value Approximation. *arXiv preprint arXiv:1306.4265* (2013).
- [30] Irwin Mann and Lloyd S Shapley. 1960. Values of Large Games, IV: Evaluating the Electoral College by Monte Carlo Techniques. (1960).
- [31] Dragos Margineantu and Thomas Dietterich. 1997. Pruning Adaptive Boosting. In *ICML*, Vol. 97. Citeseer, 211–218.
- [32] Gonzalo Martínez-Muñoz and Alberto Suárez. 2006. Pruning in Ordered Bagging Ensembles. In *Proceedings of the 23rd International Conference on Machine Learning*. 609–616.
- [33] Fatiha Mokdad, Djamel Bouchaffra, Nabil Zerrouki, and Azzedine Touazi. 2015. Determination of an Optimal Feature Selection Method Based on Maximum Shapley Value. In *2015 15th International Conference on Intelligent Systems Design and Applications (ISDA)*. IEEE, 116–121.
- [34] Annamalai Narayanan, Mahinthan Chandramohan, Rajasekar Venkatesan, Lihui Chen, and Yang Liu. 2017. Graph2Vec: Learning Distributed Representations of Graphs. (2017).
- [35] Ludwig Johann Neumann, Oskar Morgenstern, et al. 1947. *Theory of Games and Economic Behavior*. Vol. 60. Princeton University Press Princeton.
- [36] Martin J Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT press.
- [37] Guillermo Owen. 1972. Multilinear Extensions of Games. *Management Science* 18, 5-part-2 (1972), 64–79.
- [38] Miklós Pintér. 2011. Regression Games. *Annals of Operations Research* 186, 1 (2011), 263–274.
- [39] Benedek Rozemberczki, Oliver Kiss, and Rik Sarkar. 2020. Karate Club: An API Oriented Open-source Python Framework for Unsupervised Learning on Graphs. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*. ACM, 3125–3132.
- [40] Benedek Rozemberczki and Rik Sarkar. 2020. Characteristic Functions on Graphs: Birds of a Feather, from Statistical Descriptors to Parametric Models. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*. ACM, 1325–1334.
- [41] Lloyd S Shapley. 1953. A Value for n-Person Games. *Contributions to the Theory of Games* 2, 28 (1953), 307–317.
- [42] Julian Stier, Gabriele Gianini, Michael Granitzer, and Konstantin Ziegler. 2018. Analysing Neural Network Topologies: A Game Theoretic Approach. *Procedia Computer Science* 126 (2018), 234–243.
- [43] Erik Strumbelj and Igor Kononenko. 2010. An Efficient Explanation of Individual Classifications Using Game Theory. *The Journal of Machine Learning Research* 11 (2010), 1–18.
- [44] Xin Sun, Yanheng Liu, Jin Li, Jianqi Zhu, Xuejie Liu, and Huiling Chen. 2012. Using Cooperative Game Theory to Optimize the Feature Selection Problem. *Neurocomputing* 97 (2012), 86–93.
- [45] Mukund Sundararajan and Amir Najmi. 2020. The Many Shapley Values for Model Explanation. In *International Conference on Machine Learning*. PMLR, 9269–9278.
- [46] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, Alexander Bronstein, and Emmanuel Müller. 2018. NetLSD: Hearing the Shape of a Graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2347–2356.
- [47] Kagan Tumer and Joydeep Ghosh. 1996. Error Correlation and Error Reduction in Ensemble Classifiers. *Connection Science* 8, 3-4 (1996), 385–404.
- [48] Saurabh Verma and Zhi-Li Zhang. 2017. Hunt for the Unique, Stable, Sparse and Fast Feature Learning on Graphs. In *Advances in Neural Information Processing Systems*. 88–98.
- [49] Abraham Itzhak Weinberg and Mark Last. 2019. Selecting a Representative Decision Tree from an Ensemble of Decision-Tree Models for Fast Big Data Classification. *Journal of Big Data* 6, 1 (2019), 1–17.
- [50] Brian Williamson and Jean Feng. 2020. Efficient Nonparametric Statistical Inference on Population Feature Importance Using Shapley Values. In *International Conference on Machine Learning*. PMLR, 10282–10291.
- [51] Li Zhang and Wei-Da Zhou. 2011. Sparse Ensembles Using Weighted Combination Methods Based on Linear Programming. *Pattern Recognition* 44, 1 (2011), 97–106.
- [52] Yi Zhang, Samuel Burer, W Nick Street, Kristin P Bennett, and Emilio Parrado-Hernández. 2006. Ensemble Pruning Via Semi-definite Programming. *Journal of Machine Learning Research* 7, 7 (2006).
- [53] Gilad Zlotkin and Jeffrey S. Rosenschein. 1994. Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains. In *AAAI*. 432–437.