Data Valuation for Machine Learning

XIAO TIAN, National University of Singapore, Singapore

In the era of data explosion, machine learning becomes increasingly prevalent. Model owners train their models with huge amount of data from various data sources. With the development of data protection and regulation policies, data valuation becomes essential for model owners to understand the contribution of each data source to their models and how much they should compensate the data owners. As a result, accurate and efficient data valuation techniques are required in many use cases such as model interpolation, collaborative learning and domain adaption. In this survey paper, we classify current research on data valuation into 3 categories based on their methodologies: statistical methods, game-theoretic methods and meta learning methods. We discuss and evaluate the key approach used by each method in order to summarize the current state of research and provide possible directions for future works.

CCS Concepts: • **Computing methodologies → Machine learning**.

Additional Key Words and Phrases: data valuation, game theory, meta learning, Shapley value, statistics

ACM Reference Format:
Xiao Tian. 2022. Data Valuation for Machine Learning. *J. ACM* 0, 0, Article 0 (October 2022), 8 pages. https://doi.org/0000000.0000000

1 INTRODUCTION

Over the years, machine learning (ML) has been a prevalent technology to solve data-driven problems such as natural language processing, image recognition and personalization. ML is the process of training a model to learn the governing concept behind a given dataset so that it can deal with further tasks including prediction and classification. In recent ages, information sharing has become increasingly important for model owners to obtain enough data to train their models. There are two key questions that arise from the sharing process: *Who owns the data?* and *How good are the data?*. The first question asks about the *data owners*, who need to be compensated for their data; the second question asks about the *data quality*, which directly determines the final model performance. Both questions are related to the *value* of data, with which we can fairly compensate the data owners based on the impact of their data on the model performance. Therefore, it is important for us to develop suitable data valuation methods in order to accurately measure the value of data.

Classical data valuation methods are built on statistical studies, and are commonly used nowadays. There are also modern data valuation methods that provide more accurate data valuation based on advancement in other fields such as game theory and deep learning. All these methods make various assumptions or work with certain models and are hence suitable for certain use cases. Below are the major use cases of data valuation.

Author's address: Xiao Tian, National University of Singapore, 13 Computing Drive, Singapore, 117417.

0004-5411/2022/10-ART0 $15.00
https://doi.org/0000000.0000000

Model Interpretation. Data valuation explains how each data point contributes to the final model. Since higher value implies higher contribution to the model performance, we can improve the model performance by removing the group of data with lowest values and retraining the model using the rest of data with their values as weights. Recent works have empirically verified its possibility [15, 38, 46].

Collaborative ML. Collaborative ML involves the collaboration of multiple parties through data sharing, such as federated learning. After the model is trained, rewards such as monetary profits or part of the model itself should be fairly allocated among each participants. Recent works have developed various reward allocation schemes using data valuation [37, 39, 44, 47].

Domain Adaption. The distribution of data in the training set might be different from those in the validation and testing sets. This may cause existing ML methods to fail [12, 16]. Hence, we can value the data based on both their contributions to model performance and the suitable distribution so that we can resample the useful data based on their values [28, 31, 46].

In this survey paper, we summarised 3 major categories of data valuation methods based on their methodologies and discussed their pros, cons and suitable use cases.

2 METHODOLOGIES

2.1 Statistical Methods

Statisticians are interested in *robust* statistical methods so that the predicted model parameters are not largely affected by outliers. To quantitatively identify outliers, several methods to represent the influence of each data point to the model parameter have been developed, such as Cook's Distance [8] and Influence function [17].

2.1.1 Cook's Distance. In regression analysis, the *influence* of a data point represents the effect of its deletion on the regression line. In 1977, Cook [8] proposed Cook's distance which quantifies the influence of a data point by calculating the Euclidean distance between the prior and posterior model when the data point is deleted. This distance is commonly used in linear regression problems with ordinary least-squares solutions.

Though being well established and relatively scalable, the use of Cook's distance is limited to Linear Regression models in the area of ML. Even for Linear Regression models, recent research has shown that Cook's distance is outperformed by other state-of-the-art (SOTA) data valuation techniques for both high value and low value observations [15, 20].

2.1.2 Influence Function. Hampel [17] discovered the Influence curve to quantify the influence of each data point in a statistics model and Cook and Weisberg [9] developed the Influence function to quantify influence empirically. Koh and Liang [21] applied the first-order Taylor expansion of the Influence Function to simple ML models like logistic regression and showed that it can give accurate valuation of data even with a relatively large perturbation to the model. Basu *et al.* [4] argued that the change in model parameters can be large when a group of data points is removed and proposed the use of second-order Taylor expansion for group data removal.

However, Ghorbani *et al.* [13] has shown that the Influence Function used in neural networks gives vastly different valuation of data with systematic perturbation and adversarial attack. Basu *et al.* [3] also showed that the Influence function is fragile when used in deep learning models. Moreover, Koh *et al.* [22] argued that the Influence Function becomes less accurate in measuring the value of groups of data points.

Despite the robustness issues, recent research has extended the use of Influence function from directly valuing data to being used in the approximation of other data valuation techniques such as Cook's distance [43] and the Shapley value [19].

2.1.3 Leave-One-Out. The Leave-One-Out (LOO) method is another popular method in data valuation. The LOO score of a data source i is defined as the difference in the model performance with and without data source i in the dataset D. Mathematically,

$$LOO(i) = v(D) - v(D \setminus \{i\}), \tag{1}$$

where v denotes the evaluation metrics. The LOO score is a generalisation of the idea of Cook's Distance so that it can be used in any form of ML models besides Linear Regression. From Equation (1), the computation of LOO score requires $|D|$ number of model evaluations, which is relatively costly with large-scale datasets. Wang *et al.* [41] provided an efficient approximation of the LOO score using the Influence function and gradient of the loss function. However, this method is still costly when the Hessian of the loss function is hard to approximate [46].

Although LOO is a commonly used technique in data valuation, it is outperformed by most SOTA techniques and is usually used as the benchmark algorithm in related research.

2.2 Game-theoretic Methods

This group of data valuation methods is based on existing solutions to n-person cooperative games, where each data source is viewed as a player in the game. A finite cooperative game is uniquely defined by a set function $v : 2^N \longrightarrow \mathbb{R}$ with $v(\emptyset) = 0$, where N denotes the finite set of players. This function measures the utility of each subset of players, i.e. *coalition* of players, and in our case, the evaluation metrics. The goal is to find a function ϕ such that $\phi v : N \longrightarrow \mathbb{R}$ fairly measures the contribution of player i in N. Dubey *et al.* [11] introduced 4 axioms to define fairness: *Linearity*, *Symmetry*, *Monotonicity* and *Projection* and discovered a collection of functions, namely semivalues, that satisfy these axioms. The mathematical expression of a semivalue is as follows:

$$(\phi v)(i) = \sum_{S \subseteq N \setminus i} w_s [v(S \cup \{i\}) - v(S)]. \tag{2}$$

The term $v(S \cup \{i\}) - v(S)$ represents the *marginal contribution* of player i to coalition s, and w_s represents the weight assigned to each marginal contribution to coalitions of size s. Note that w needs to satisfy $\sum_{s=0}^{n-1} \binom{n-1}{s} w_s = 1$, where n denotes the number of players in N. The notations used in Equation (2) will be consistent throughout the rest of this paper.

Most data valuation methods based on cooperative game theory (CGT) focus on finding the most appropriate semivalue under certain problem setting by adjusting the value of w. Meanwhile, the marginal contributions are extremely costly to compute, hence another major research direction is to efficiently approximate the semivalues. Table 1 summarises the current algorithms to efficiently approximate the actual values.

2.2.1 Data Shapley. Ghorbani *et al.* [15] and Jia *et al.* [19] extended the definition of fairness to the equitable valuation of data sources and formulated their data valuation function based on the Shapley Value (SV). SV is a type of semivalue where $w_s = \frac{1}{n}\binom{n-1}{s}^{-1}$ in Equation (2) [33]. Followed from SV, the Data Shapley value is defined as follows:

$$\phi_i = C \sum_{S \subseteq D \setminus \{i\}} \binom{n-1}{|S|}^{-1} [v(S \cup \{i\}) - v(S)], \tag{3}$$

Table 1. Computational cost of various game-theoretic data valuation algorithms applied on a dataset of size n. LOO method is also included for reference.

Method	Number of Model Evaluation Needed	Approximation Error
LOO	$O(n)$	0
Data Shapley	$O(2^n)$	0
TMC-Shapley [15]	$\approx O(mn)$	$\approx O(\sqrt{\frac{r}{m}})$
G-Shapley [15]	$\approx$ TMC-Shapley	-
Permutation Sampling [19]	$O(n^2 \log n)$	(ϵ, δ)
Group Testing [19]	$O\left(n(\log n)^2\right)$	(ϵ, δ)
$\mathcal{D}$-Shapley [14]	$O\left(\frac{\log(\frac{n}{\delta})}{\epsilon^2}\right)$	(ϵ, δ)
MSR-Banzhaf [40]	$O\left(\frac{n}{\epsilon^2} \log \frac{n}{\delta}\right)$	(ϵ, δ)
Least Core [45]	$O\left(\frac{\left(\frac{\max_S v(S)}{\min_{S\neq\emptyset} v(S)}\right)^2 \left(\log n + \log \frac{1}{\delta}\right)}{\epsilon^2 \delta^2}\right)$	(ϵ, δ)

where C denotes an arbitrary constant, D denotes the dataset and other notations are same as Equation (2).

In order to efficiently approximate the Data Shapley value, Ghorbani *et al.* proposed the Truncated Monte Carlo Shapley (TMC-Shapley) algorithm and the Gradient Shapley (G-Shapley) algorithm. Both algorithms outperform the LOO method empirically. Jia *et al.* proposed a different algorithm exploiting group testing and sparsity of values, which is proven to guarantee the approximation error. They also proposed two practical algorithms based on stable learning algorithms and influence functions, for which the approximation error is not proven to be guaranteed.

Approximation of Data Shapley value in specific ML models has also been studied. Jia *et al.* [18] proposed an algorithm for k-Nearest Neighbour classifier that runs in sublinear time. Ancona *et al.* [1] proposed the DASP algorithm which requires a polynomial number of network evaluations.

2.2.2 Distributional Shapley. Ghorbani *et al.* [14] argued that the valuation made by Data Shapley depends on the actual dataset, which neither accounts for the original statistical distribution of data nor provides insights on data points outside the given dataset. Therefore, they proposed Distributional Shapley, which values the data considering the underlying statistical distribution.

One key advantage of the Distributional Shapley framework is that the value of each data source is stable when different datasets are sampled from the original pool of data. Moreover, unlike Data Shapley where the value of each data source severely depends on other data sources, the Distributional Shapley value of each data source is private to the data source itself, which provides privacy for data owners. On the other hand, the specific algorithm for the Distributional Shapley framework depends on the models and value functions, reducing its universality.

2.2.3 Beta Shapley. Kwon and Zou [24] proved that the choice of w_s in Data Shapley is suboptimal to reflect the influence of individual data. Since the marginal contribution of each data source to smaller coalitions tends to have larger signal-to-noise ratio, it is reasonable to assign a larger weight to these marginal contributions. Kwon and Zou defined the Beta Shapley value as follows:

$$\phi_i = \sum_{S \subseteq D \setminus \{i\}} \frac{Beta(|S| + \beta, n - |S| - 1 + \alpha)}{Beta(\alpha, \beta)} [v(S \cup \{i\}) - v(S)] \tag{4}$$

The use of Beta distribution brings more freedom to adjust the weights. For example, when $\alpha = 1$, $\beta = 1$, Equation (4) becomes Data Shapley. When $\alpha \geq \beta = 1$, a larger weight is assigned to marginal contributions to smaller coalitions. The authors empirically discovered that Beta Shapley with $Beta(16, 1)$ outperforms Data Shapley and other SOTA methods.

The authors employed Monte Carlo approximation similar as what Ghorbani *et al.* did in Data Shapley [15]. Such algorithm is not efficient enough for large-scale datasets in practice.

2.2.4 Data Banzhaf. Wang and Jia [40] defined *robustness* of data valuation as the amount of perturbation to the model performance scores so that the order of data values does not change. They proved that the Banzhaf index [2] uniquely gives the maximal robustness and defined Data Banzhaf value as follows:

$$\phi_i = \sum_{S \subseteq D \setminus \{i\}} \frac{1}{2^{n-1}} [v(S \cup \{i\}) - v(S)] \tag{5}$$

In order to efficiently approximate the Data Banzhaf value, the authors used the Maximum Sample Reuse algorithm. This algorithm outperforms Monte Carlo method and is uniquely applicable to valuation methods based on Banzhaf value, which marked a progress towards the empirical use of semivalue-based data valuation methods. One possible drawback is that it put too many weights on coalitions of size close to $\frac{n}{2}$, which might be an over-representation of such terms.

2.2.5 Variational Index. The last two methods based on CGT are not derived from semivalues. Bian *et al.* [5] introduced an energy-based [25] treatment for cooperative games, namely the Variational Index. In order to remove the correlation between players to measure their individual values, we need to decouple their interactions by minimising the best conceivable decoupling distance. This minimisation problem can be solved using multilinear extension [6, 30]. The authors proved that the Variational Index value satisfies a set of desirable axioms.

As a different game-theoretic approach, the Variational Index method empirically outperforms other SOTA data valuation methods sometimes, and always achieves the lowest decoupling error.

2.2.6 Least Core. Another equally important concept to SV in CGT is the *least core* [27], which represents the case where the maximum deficit of any coalition is minimised. Yan *et al.* [45] proposed a way to compute the least core by solve the linear programming:

$$\min \quad e \quad s.t. \quad \sum_{i \in N} \phi_i = v(N) \text{ and } \sum_{i \in S} \phi_i + e \geq v(S), \forall S \subset N \tag{6}$$

The authors used Monte Carlo sampling as an approximation and theoretically guaranteed the time complexity and approximation error. Their experiments showed that the Least Core method outperforms TMC-Shapley and Group Testing Shapley in several data removal tasks. Although this method also faces complexity issues, it provides a new perspective to value data.

2.3 Meta Learning Methods

There are past research on using meta learning to compute adaptive weights in robust learning [7, 26, 34, 35]. However, it was only until recently that meta learning was applied to data valuation. Unlike game-theoretic approach, meta learning methods do not make assumptions on the problem setting.

2.3.1 Data Valuation with Reinforcement Learning. Yoon *et al.* [46] adapted the meta learning methodology from robust learning to data valuation. Their proposed model, Data Valuation with Reinforcement Learning (DVRL), aims to learn the predictor and the data valuation function together. The predictor is trained through normal stochastic gradient descent method whereas the non-differentiable data valuation function is trained using the REINFORCE algorithm [42].

The computational cost of DVRL is much lower than that of statistical or game-theoretic methods. It is not related to the size of dataset but rather the number of iterations and training complexity per iteration. Actually, the training time of DVRL is approximately twice of the conventional training, which is much more scalable. Meanwhile, the performance of DVRL is not worse than other SOTA methods such as Data Shapley. DVRL outperforms Data Shapley in domain adaption and robust learning tasks and performs similarly for other data valuation tasks.

One limitation of DVRL is that it does not entertain the various equitable valuation axioms. This may restrict the use of DVRL in the case where equitability is important, such as reward allocation.

3 DISCUSSIONS AND FUTURE WORK

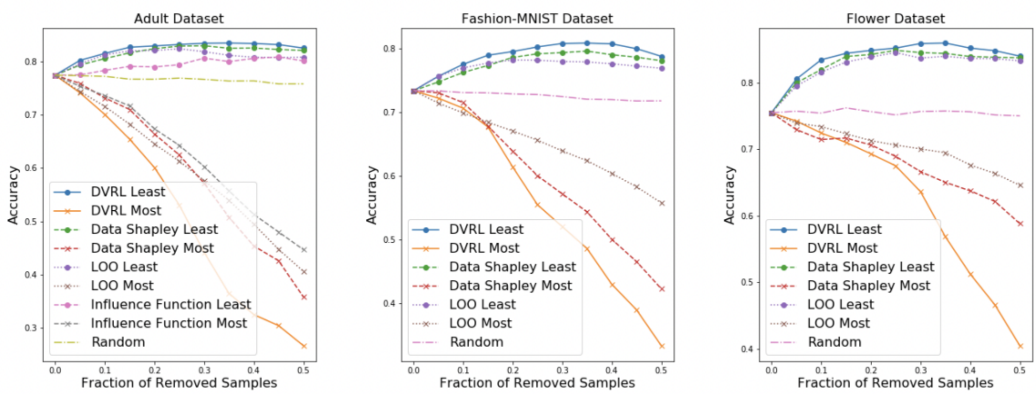

Fig. 1. Comparisons of performances of various data valuation techniques in data removal tasks [46]. Predicting accuracy changes when the largest (marked as •) and smallest (marked as ×) value samples are removed.

Figure 1 shows a comparison of data valuation methods from all categories we discussed. Generally, methods based on CGT and meta learning outperform traditional methods.

As the problem settings in ML are not exactly the same as those in CGT, the desirable axioms which the data valuation function needs to satisfy have always been under discussion. The same set of axioms may not be optimal for all problem settings [5, 32]. Besides the original axioms of semivalues, Covert *et al.* [10], Ridaoui *et al.* [32] and Sim *et al.* [36] have provided different sets of axioms that are considered favourable in ML. This drives the need for different expression of semivalue-based methods in different problem settings, which is a possible direction for future work. For example, Ohrimenki *et al.* [29] defined *Replication Robustness* axiom to prevent unethical replication of data for more rewards and developed Robust Shapley valuation function.

Another challenge faced by most statistics and game-theoretic methods is their computational cost. Most of them requires as least $O(n)$ number of model evaluations, where n represents the size of dataset, and the cost is typically higher for game-theoretic methods. Such computational cost is too high for training of complex models such as deep neural networks. Moreover, most of the current approximation algorithms, as shown in Table 1, do not provide a theoretically bounded approximation error. This leaves concern for the empirical use of such algorithms. Therefore, further work can explore more efficient approximation to such value functions.

Although game-theoretic methods seem promising and works well empirically, there are some doubts about the correct usage of such approaches. Kumar *et al.* [23] argued that both SV and the core are not suitable for non-additive games, but most ML problems are taken as additive games for granted. Further work can explore the alternatives to such concepts which can be applied to non-additive games, in order to better generalise this category of approach.

REFERENCES

[1] Marco Ancona, Cengiz Oztireli, and Markus Gross. 2019. Explaining deep neural networks with a polynomial time algorithm for shapley value approximation. In *International Conference on Machine Learning*. PMLR, 272–281.

[2] John F Banzhaf III. 1964. Weighted voting doesn't work: A mathematical analysis. *Rutgers L. Rev.* 19 (1964), 317.

[3] Samyadeep Basu, Philip Pope, and Soheil Feizi. 2020. Influence functions in deep learning are fragile. *arXiv preprint arXiv:2006.14651* (2020).

[4] Samyadeep Basu, Xuchen You, and Soheil Feizi. 2020. On second-order group influence functions for black-box predictions. In *International Conference on Machine Learning*. PMLR, 715–724.

[5] Yatao Bian, Yu Rong, Tingyang Xu, Jiaxiang Wu, Andreas Krause, and Junzhou Huang. 2021. Energy-Based Learning for Cooperative Games, with Applications to Valuation Problems in Machine Learning. *arXiv preprint arXiv:2106.02938* (2021).

[6] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. 2007. Maximizing a submodular set function subject to a matroid constraint. In *International Conference on Integer Programming and Combinatorial Optimization*. Springer, 182–196.

[7] Sungjoon Choi, Sanghoon Hong, Kyungjae Lee, and Sungbin Lim. 2018. Choicenet: Robust learning by revealing output correlations. (2018).

[8] R Dennis Cook. 1977. Detection of influential observation in linear regression. *Technometrics* 19, 1 (1977), 15–18.

[9] R Dennis Cook and Sanford Weisberg. 1980. Characterizations of an empirical influence function for detecting influential cases in regression. *Technometrics* 22, 4 (1980), 495–508.

[10] Ian Covert, Scott M Lundberg, and Su-In Lee. 2021. Explaining by Removing: A Unified Framework for Model Explanation. *J. Mach. Learn. Res.* 22 (2021), 209–1.

[11] Pradeep Dubey, Abraham Neyman, and Robert James Weber. 1981. Value theory without efficiency. *Mathematics of Operations Research* 6, 1 (1981), 122–128.

[12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research* 17, 1 (2016), 2096–2030.

[13] Amirata Ghorbani, Abubakar Abid, and James Zou. 2019. Interpretation of neural networks is fragile. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 3681–3688.

[14] Amirata Ghorbani, Michael Kim, and James Zou. 2020. A distributional framework for data valuation. In *International Conference on Machine Learning*. PMLR, 3535–3544.

[15] Amirata Ghorbani and James Zou. 2019. Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*. PMLR, 2242–2251.

[16] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*.

[17] Frank R. Hampel. 1974. The Influence Curve and Its Role in Robust Estimation. *J. Amer. Statist. Assoc.* 69, 346 (1974), 383–393. http://www.jstor.org/stable/2285666

[18] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nezihe Merve Gurel, Bo Li, Ce Zhang, Costas J Spanos, and Dawn Song. 2019. Efficient task-specific data valuation for nearest neighbor algorithms. *arXiv preprint arXiv:1908.08619* (2019).

[19] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. 2019. Towards efficient data valuation based on the shapley value. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 1167–1176.

[20] Mattias Jönsson. 2020. On Valuation of Observations in Linear Regression Models. (2020).

[21] Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*. PMLR, 1885–1894.

[22] Pang Wei W Koh, Kai-Siang Ang, Hubert Teo, and Percy S Liang. 2019. On the accuracy of influence functions for measuring group effects. *Advances in neural information processing systems* 32 (2019).

[23] I Elizabeth Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, and Sorelle Friedler. 2020. Problems with Shapley-value-based explanations as feature importance measures. In *International Conference on Machine Learning*. PMLR, 5491–5500.

[24] Yongchan Kwon and James Zou. 2021. Beta shapley: a unified and noise-reduced data valuation framework for machine learning. *arXiv preprint arXiv:2110.14049* (2021).

[25] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. 2006. A tutorial on energy-based learning. *Predicting structured data* 1, 0 (2006).

[26] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. 2019. Learning to learn from noisy labeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5051–5059.

[27] Michael Maschler, Bezalel Peleg, and Lloyd S Shapley. 1979. Geometric properties of the kernel, nucleolus, and related solution concepts. *Mathematics of operations research* 4, 4 (1979), 303–338.
[28] Jiquan Ngiam, Daiyi Peng, Vijay Vasudevan, Simon Kornblith, Quoc V Le, and Ruoming Pang. 2018. Domain adaptive transfer learning with specialist models. *arXiv preprint arXiv:1811.07056* (2018).
[29] Olga Ohrimenko, Shruti Tople, and Sebastian Tschiatschek. 2019. Collaborative machine learning markets with data-replication-robust payments. *arXiv preprint arXiv:1911.09052* (2019).
[30] Guillermo Owen. 1972. Multilinear extensions of games. *Management Science* 18, 5-part-2 (1972), 64–79.
[31] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. In *International conference on machine learning*. PMLR, 4334–4343.
[32] Mustapha Ridaoui, Michel Grabisch, and Christophe Labreuche. 2018. An axiomatisation of the Banzhaf value and interaction index for multichoice games. In *International Conference on Modeling Decisions for Artificial Intelligence*. Springer, 143–155.
[33] Lloyd S Shapley. 1997. A value for n-person games. *Classics in game theory* 69 (1997).
[34] Yanyao Shen and Sujay Sanghavi. 2019. Learning with bad training data via iterative trimmed loss minimization. In *International Conference on Machine Learning*. PMLR, 5739–5748.
[35] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. *Advances in neural information processing systems* 32 (2019).
[36] Rachael Hwee Ling Sim, Xinyi Xu, and Bryan Kian Hsiang Low. 2022. Data valuation in machine learning:"ingredients", strategies, and open challenges. In *Proc. IJCAI*.
[37] Rachael Hwee Ling Sim, Yehong Zhang, Mun Choon Chan, and Bryan Kian Hsiang Low. 2020. Collaborative machine learning with incentive-aware model rewards. In *International Conference on Machine Learning*. PMLR, 8927–8936.
[38] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. 2020. Dataset cartography: Mapping and diagnosing datasets with training dynamics. *arXiv preprint arXiv:2009.10795* (2020).
[39] Sebastian Shenghong Tay, Xinyi Xu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. 2022. Incentivizing collaboration in machine learning via synthetic data rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 9448–9456.
[40] Tianhao Wang and Ruoxi Jia. 2022. Data Banzhaf: A Data Valuation Framework with Maximal Robustness to Learning Stochasticity. *arXiv preprint arXiv:2205.15466* (2022).
[41] Zifeng Wang, Hong Zhu, Zhenhua Dong, Xiuqiang He, and Shao-Lun Huang. 2020. Less is better: Unweighted data subsampling via influence function. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 6340–6347.
[42] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3 (1992), 229–256.
[43] Mike Wojnowicz, Ben Cruz, Xuan Zhao, Brian Wallace, Matt Wolff, Jay Luan, and Caleb Crable. 2016. "Influence sketching": Finding influential samples in large-scale regressions. In *2016 IEEE International Conference on Big Data (Big Data)*. IEEE, 3601–3612.
[44] Xinyi Xu, Lingjuan Lyu, Xingjun Ma, Chenglin Miao, Chuan Sheng Foo, and Bryan Kian Hsiang Low. 2021. Gradient driven rewards to guarantee fairness in collaborative machine learning. *Advances in Neural Information Processing Systems* 34 (2021), 16104–16117.
[45] Tom Yan and Ariel D Procaccia. 2021. If you like shapley then you'll love the core. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5751–5759.
[46] Jinsung Yoon, Sercan Arik, and Tomas Pfister. 2020. Data valuation using reinforcement learning. In *International Conference on Machine Learning*. PMLR, 10842–10851.
[47] Jingwen Zhang, Yuezhou Wu, and Rong Pan. 2021. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In *Proceedings of the Web Conference 2021*. 947–956.