Collaborative Machine Learning with Incentive-Aware Model Rewards

Rachael Hwee Ling Sim [1] **Yehong Zhang** [1] **Mun Choon Chan** [1] **Bryan Kian Hsiang Low** [1]

Abstract

Collaborative *machine learning* (ML) is an appealing paradigm to build high-quality ML models by training on the aggregated data from many parties. However, these parties are only willing to share their data when given enough incentives, such as a guaranteed fair reward based on their contributions. This motivates the need for measuring a party's contribution and designing an incentive-aware reward scheme accordingly. This paper proposes to value a party's reward based on Shapley value and information gain on model parameters given its data. Subsequently, we give each party a model as a reward. To formally incentivize the collaboration, we define some desirable properties (e.g., fairness and stability) which are inspired by cooperative game theory but adapted for our model reward that is uniquely *freely replicable*. Then, we propose a novel model reward scheme to satisfy fairness and trade off between the desirable properties via an adjustable parameter. The value of each party's model reward determined by our scheme is attained by injecting Gaussian noise to the aggregated training data with an optimized noise variance. We empirically demonstrate interesting properties of our scheme and evaluate its performance using synthetic and real-world datasets.

1. Introduction

Collaborative *machine learning* (ML) is an appealing paradigm to build high-quality ML models. While an individual party may have limited data, it is possible to build improved, high-quality ML models by training on the aggregated data from many parties. For example, in healthcare, a hospital or healthcare firm whose data diversity and quantity are limited due to its small patient base can draw on data from other hospitals and firms to improve the prediction of some disease progression (e.g., diabetes) (Center for Open Data Enterprise, 2019). This collaboration can be encouraged by a government agency, such as the National Institute of Health in the United States. In precision agriculture, a farmer with limited land area and sensors can combine his collected data with the other farmers to improve the modeling of the effect of various influences (e.g., weather, pest) on his crop yield (Claver, 2019). Such data sharing also benefits other application domains, including real estate in which a property agency can pool together its limited transactional data with that of the other agencies to improve the prediction of property prices (Conway, 2018).

However, any party would have incurred some nontrivial cost to collect its data. So, they would not altruistically donate their data and risk losing their competitive edge. These parties will be motivated to share their data when given enough incentives, such as a guaranteed benefit from the collaboration and a *fair* higher reward from contributing more valuable data. To this end, we propose to value each party's contributed data and design an incentive-aware reward scheme to give each party a separate ML model as a reward (in short, *model reward*) accordingly. We use only model rewards and exclude monetary compensations as (a) in the above-mentioned applications, every party such as a hospital is mainly interested in improving model quality for unlimited future test predictions; (b) there may not be a feasible and clear source of monetary revenue to compensate participants (e.g., due to restrictions, the government may not be able to pay firms using tax revenues); and (c) if parties have to pay to participate in the collaboration, then the total payment is debatable and many may lack funding while still having valuable data.

How then should we value a party's data and its effect on model quality? To answer this first question, we propose a valuation method based on the *informativeness* of data. In particular, more informative data induces a greater reduction in the uncertainty of the model parameters, hence improving model quality. In contrast, existing data valuation methods (Ghorbani & Zou, 2019; Ghorbani et al., 2020; Jia et al., 2019a;b; Yoon et al., 2020) measure model quality via its validation accuracy, which calls for a tedious or even impossible process of needing all parties to agree on a common validation dataset. Inaccurate valuation can also arise due to

[1]Department of Computer Science, National University of Singapore, Republic of Singapore. Correspondence to: Bryan Kian Hsiang Low <lowkh@comp.nus.edu.sg>.

Proceedings of the 37^{th} International Conference on Machine Learning, Online, PMLR 119, 2020.

how a party's test queries, which are likely to change over time, differ from the validation dataset. Our data valuation method does not make any assumption about the current or future distribution of test queries.

Next, *how should we design a reward scheme to decide the values of model rewards for incentivizing a collaboration?* Intuitively, a party will be motivated to collaborate if it can receive a better ML model than others who have contributed less valuable data, and than what it can build alone or get from collaborating separately with some parties. Also, the parties often like to maximize the total benefit from the collaboration. These incentives appear related to solution concepts (*fairness*, *individual rationality*, *stability*, and *group welfare*) from *cooperative game theory* (CGT), respectively. However, as CGT assumptions are restrictive for the uniquely *freely replicable*[1] nature of our model reward, these CGT concepts have to be adapted for defining incentives for our model reward. We then design a novel model reward scheme to provide these incentives.

Finally, *how can we realize the values of the model rewards decided by our scheme?* How should we modify the model rewards or the data used to train them? An obvious approach to control the values of the model rewards is to select and only train on subsets of the aggregated data. However, this requires considering an exponential number of *discrete* subsets of data, which is intractable for large datasets such as medical records. We avoid this issue by injecting noise into the aggregated data from multiple parties instead. The value of each party's model reward can then be realized by simply optimizing the *continuous* noise variance parameter.

The specific contributions of our work here include:

- Proposing a data valuation method using the *information gain* (IG) on model parameters given the data (Sec. 3);
- Defining new conditions for incentives (i.e., *Shapley fairness*, *stability*, *individual rationality*, and *group welfare*) that are suitable for the *freely replicable* nature of our model reward (Sec. 4.1). As these incentives cannot be provided all at the same time, we design a novel model reward scheme with an adjustable parameter to trade off between them while maintaining fairness (Sec. 4.2);
- Injecting Gaussian noise into the aggregated data from multiple parties and optimizing the noise variance parameter for realizing the values of the model rewards decided by our scheme (Sec. 4.3); and
- Demonstrating interesting properties of our model reward scheme empirically and evaluating its performance with synthetic and real-world datasets (Sec. 5).

To the best of our knowledge, our work here is the first to propose a collaborative ML scheme that formally considers incentives beyond fairness and relies solely on model rewards to realize them. Existing works (Jia et al., 2019a;b; Ohrimenko et al., 2019) have only looked at fairness and have to resort to monetary compensations if considered.

[1]Data and model reward, like digital goods, can be replicated at zero marginal cost and given to more parties.

2. Problem Formulation

In our problem setting, we consider n honest and non-malicious parties, each of whom owns some data and assume the availability of a trusted *central party*[2] who aggregates data from these parties, measures the value of their data, and distributes a resulting trained ML model to each party. We first introduce the notations and terminologies used in this work: Let $N \triangleq \{1, \ldots, n\}$ denote a set of n parties. Any subset $C \subseteq N$ is called a *coalition* of parties. The *grand coalition* is the set N of all parties. Parties will team up and partition themselves into a *coalition structure* CS. Formally, CS is a set of coalitions such that $\bigcup_{C \in CS} C = N$ and $C \cap C' = \emptyset$ for any $C, C' \in CS$ and $C \neq C'$. The data of party $i \in N$ is represented by $D_i \triangleq (\mathbf{X}_i, \mathbf{y}_i)$ where $\mathbf{X}_i$ and $\mathbf{y}_i$ are the input matrix and output vector, respectively. Let v_C denote the *value of the (aggregated) data* $D_C \triangleq \{D_i\}_{i \in C}$ of any coalition $C \subseteq N$. We use v_i to represent $v_{\{i\}}$ to ease notation. For each party $i \in N$, r_i denotes the *value of its received model reward*.

The objective is to design a collaborative ML scheme for the central party to decide and realize the values $(r_i)_{i \in N}$ of model rewards distributed to parties $1, \ldots, n$. The scheme should satisfy certain incentives (e.g., fairness and stability) to encourage the collaboration. Some works (Ghorbani & Zou, 2019; Jia et al., 2019a;b) have considered similar problems and can fairly *partition* the (monetary) value v_N of the entire aggregated data into r_i for $i \in N$. They achieved this using results from *cooperative game theory* (CGT). Our problem, however, differs and cannot be addressed directly using CGT. This is due to the *freely replicable* nature of our *model reward* – the total value $\sum_{i \in N} r_i$ of received model rewards over all parties $i \in N$ can exceed v_N.

We will next show how to assess the value of data (Sec. 3) and, more importantly, how to design the reward scheme to decide the values $(r_i)_{i \in N}$ of model rewards accordingly and realize these values for achieving our incentive-aware objective (Sec. 4).

3. Data Valuation with Information Gain

A set D_C of data is considered more valuable (i.e., higher v_C) if it can be used to train a higher-quality (hence more valuable) ML model. Existing data valuation methods

[2]In reality, such a central party can be found in established data sharing platforms like Ocean Protocol (Ocean Protocol Foundation, 2019) and Data Republic (`https://datarepublic.com`).

(Ghorbani & Zou, 2019; Ghorbani et al., 2020; Jia et al., 2019b; Yoon et al., 2020) measure the quality of a trained ML model via its validation accuracy. However, these methods require the central party to carefully select a validation dataset that all parties must agree on. This is often a tedious or even impossible process, especially if every party's test queries, which are likely to change over time, differ from the validation dataset.[3] For example, two private hospitals $\mathcal{H}_1$ and $\mathcal{H}_2$ aggregate their data for diabetes prediction. $\mathcal{H}_1$ and $\mathcal{H}_2$ prefer accurate test predictions for female and young patients, respectively. Due to the differences in their data sizes, data qualities, and preferred test queries, it is difficult for the central party to decide the demographics of the patients in the validation dataset such that the data valuation is unbiased and accurate.

To circumvent the above-mentioned issues, our proposed data valuation method instead considers an information-theoretic measure of the quality of a trained model in terms of the reduction in uncertainty of the model parameters, denoted by vector $\boldsymbol{\theta}$, after training on data D_C. We use the prior entropy $\mathbb{H}(\boldsymbol{\theta})$ and posterior entropy $\mathbb{H}(\boldsymbol{\theta}|D_C)$ to represent the uncertainty of $\boldsymbol{\theta}$ before and after training on D_C, respectively. So, if the data D_C for $C \subseteq N$ can induce a greater reduction in the uncertainty/entropy of $\boldsymbol{\theta}$ or, equivalently, *information gain* (IG) $\mathbb{I}(\boldsymbol{\theta}; D_C)$ on $\boldsymbol{\theta}$:

$$v_C \triangleq \mathbb{I}(\boldsymbol{\theta}; D_C) = \mathbb{H}(\boldsymbol{\theta}) - \mathbb{H}(\boldsymbol{\theta}|D_C)\ , \tag{1}$$

then a higher-quality (hence more valuable) model can be trained using this more valuable/informative data D_C. IG is an appropriate data valuation method as it is often used as a surrogate measure of the test/predictive accuracy of a trained model (Krause & Guestrin, 2007; Kruschke, 2008) since the test queries are usually not known *a priori*. We empirically demonstrate such a surrogate effect in Appendix F.3. The predictive distribution of the output y^* at the test input $\mathbf{x}^*$ given data D_C is calculated by averaging over all possible model parameters $\boldsymbol{\theta}$ weighted by their posterior belief $p(\boldsymbol{\theta}|D_C)$, i.e., $p(y^*|\mathbf{x}^*, D_C) = \int p(y^*|\mathbf{x}^*, \boldsymbol{\theta})\, p(\boldsymbol{\theta}|D_C)\, \mathrm{d}\boldsymbol{\theta}$. By reducing the uncertainty in $\boldsymbol{\theta}$, we can further rule out models that are unlikely given D_C and place higher weights (i.e., by increasing $p(\boldsymbol{\theta}|D_C)$) on models that are closer to the true model parameters, thus improving the predictive accuracy for any test query in expectation. The value v_C (1) of data D_C has the following properties:

- Data of an *empty* coalition has *no* value: $v_\emptyset = 0$.
- Data of any coalition $C \subseteq N$ has *non-negative* value: $\forall C \subseteq N\ \ v_C \geq 0$.
- **Monotonicity**. Adding more parties to a coalition cannot decrease the value of its data: $\forall C \subseteq C' \subseteq N\ \ v_{C'} \geq v_C$.
- **Submodularity**. Data of any party i is less valuable to a larger coalition which has more parties and data:

$\forall i \in N\ \forall C \subseteq C' \subseteq N \backslash \{i\}\ \ v_{C' \cup \{i\}} - v_{C'} \leq v_{C \cup \{i\}} - v_C$.

The second and third properties are due to the "information never hurts" bound for entropy (Cover & Thomas, 1991). The submodular property of IG (1) assumes conditional independence of the data D_i and D_j given $\boldsymbol{\theta}$ for any $i, j \in N$ and $i \neq j$; its proof is in Appendix A.

The first two properties fulfill standard assumptions of CGT. The latter two properties will influence the design and properties of our model reward scheme in Sec. 4. For example, the monotonic property ensures that the value r_i of any party i's model reward is never negative (Sec. 4).

4. Incentive-Aware Reward Scheme with Model Rewards

Recall our problem setting in Sec. 2 that the central party will train a model for each party as a reward. The reward should be decided based on the value of each party's data relative to that of the other parties'. Let D_i^r be the data (i.e., derived from D_N) used to train party i's model reward. The value of party i's model reward is $r_i \triangleq \mathbb{I}(\boldsymbol{\theta}; D_i^r)$ according to Sec. 3. In this section, we will first present the incentives to encourage collaboration (Sec. 4.1), then describe how our incentive-aware reward scheme will satisfy them (Sec. 4.2), and finally show how to vary D_i^r to realize the values of the model rewards decided by our scheme (Sec. 4.3).

We desire a collaboration involving the *grand coalition* (i.e., $CS = \{N\}$) as it results in the largest aggregated data D_N and hence allows the highest value of model reward to be achieved, which eliminates the tedious process of deciding how the parties should team up and partition themselves. In Sec. 4.1.2, we will discuss how to incentivize the grand coalition formation and increase the total benefit from the collaboration.

4.1. Incentives

Our reward scheme has to be valid, fair to all the parties, and guarantee an improved model for each party. To achieve these, we exploit and adapt key assumptions and constraints about rewards in CGT. As has been mentioned in Sec. 2, the modifications are necessary as the model reward is uniquely freely replicable. We require the following incentive conditions to hold for the values $(r_i)_{i \in N}$ of model rewards based on the chosen coalition structure CS:

R1 **Non-negativity.** $\forall i \in N\ \ r_i \geq 0$.

R2 **Feasibility.** The model reward received by each party in any coalition $C \in CS$ cannot be more valuable

[3] In the unlikely event that all parties can agree on a common validation dataset, validation accuracy would be the preferred measure of model quality and hence data valuation method.

than the model trained on their aggregated data D_C: $\forall C \in CS \;\; \forall i \in C \;\; r_i \leq v_C$.

R3 **Weak Efficiency.** In each coalition $C \in CS$, the model reward received by at least a party $i \in C$ is as valuable as the model trained on the aggregated data D_C of C: $\forall C \in CS \;\; \exists i \in C \;\; r_i = v_C$.

R4 **Individual Rationality.** Each party should receive a model reward that is at least as valuable as the model trained on its own data: $\forall i \in N \;\; r_i \geq v_i$.

R1 and R4 are the same as the solution concepts in CGT while R2 and R3 have been adapted.[4] When $CS = \{N\}$ (i.e., grand coalition), R2 and R3 become $\forall i \in N \;\; r_i \leq v_N$ and $\exists i \in N \;\; r_i = v_N$, respectively. So, each party cannot receive a more valuable model reward than the model trained on D_N as it would involve creating data artificially.

4.1.1. Fairness

In addition, when $CS = \{N\}$ (i.e., grand coalition), to guarantee that the reward scheme is *fair* to all n parties, the values $(r_i)_{i \in N}$ of their model rewards must satisfy the following properties which are inspired by the fairness concepts in CGT (Maschler & Peleg, 1966; Shapley, 1953; Young, 1985) and have also been experimentally studied from a normative perspective (de Clippel & Rozen, 2013):

F1 **Uselessness.**[5] If including the data of party i does not improve the quality of a model trained on the aggregated data of any coalition (e.g., when $D_i = \emptyset$), then party i should receive a valueless model reward: For all $i \in N$,

$$(\forall C \subseteq N \setminus \{i\} \;\; v_{C \cup \{i\}} = v_C) \Rightarrow r_i = 0 \; .$$

F2 **Symmetry.**[5] If including the data of party i yields the same improvement as that of party j in the quality of a model trained on the aggregated data of any coalition (e.g., when $D_i = D_j$), then they should receive equally valuable model rewards: For all $i, j \in N$ s.t. $i \neq j$,

$$(\forall C \subseteq N \setminus \{i, j\} \;\; v_{C \cup \{i\}} = v_{C \cup \{j\}}) \Rightarrow r_i = r_j \; .$$

F3 **Strict Desirability (Maschler & Peleg, 1966).** If the quality of a model trained on the aggregated data of at least a coalition improves more by including the data of party i than that of party j, but the reverse is not true, then party i should receive a more valuable model reward than party j: For all $i, j \in N$ s.t. $i \neq j$,

$$(\exists B \subseteq N \setminus \{i, j\} \;\; v_{B \cup \{i\}} > v_{B \cup \{j\}}) \wedge$$
$$(\forall C \subseteq N \setminus \{i, j\} \;\; v_{C \cup \{i\}} \geq v_{C \cup \{j\}}) \Rightarrow r_i > r_j \; .$$

F4 **Strict Monotonicity.**[6] If the quality of a model trained on the aggregated data of at least a coalition containing party i improves (e.g., by including more data of party i), *ceteris paribus*, then party i should receive a more valuable model reward than before: Let $\{v_C\}_{C \in 2^N}$ and $\{v'_C\}_{C \in 2^N}$ denote any two sets of values of data over all coalitions $C \subseteq N$, and r_i and r'_i be the corresponding values of model rewards received by party i. For all $i \in N$,

$$(\exists B \subseteq N \setminus \{i\} \;\; v'_{B \cup \{i\}} > v_{B \cup \{i\}}) \wedge$$
$$(\forall C \subseteq N \setminus \{i\} \;\; v'_{C \cup \{i\}} \geq v_{C \cup \{i\}}) \wedge$$
$$(\forall A \subseteq N \setminus \{i\} \;\; v'_A = v_A) \wedge (v'_N > r_i) \Rightarrow r'_i > r_i \; .$$

We have the following incentive condition:

R5 **Fairness.** The values $(r_i)_{i \in N}$ of model rewards must satisfy F1 to F4.

Both F3 and F4 imply that marginal contribution $(v_{C \cup \{i\}} - v_C)$ matters. F4 is the only property guaranteeing that if party i adds more valuable/informative data to a coalition containing it, *ceteris paribus*, then it should receive a more valuable model reward than before. Additionally, F3 establishes a relationship between parties i and j that if their marginal contributions only differ for coalition C (i.e., $v_{C \cup \{i\}} > v_{C \cup \{j\}}$ w.l.o.g.), then party i should receive a more valuable model reward than party j.

To illustrate their significance, we consider two simpler reward schemes where we directly set the value of model reward received by every party $i \in N$ as (a) the value of its data (i.e., $r_i \triangleq v_i$) or (b) the decrease in the value of its data if it leaves the grand coalition (i.e., $r_i \triangleq v_N - v_{N \setminus \{i\}}$). Both schemes violate F3 and F4 as they ignore the values of the other parties' data: The value of model reward received by party i does not change when its marginal contribution to any non-empty coalition $C \subset N \setminus \{i\}$ increases. Intuitively, a party with a small value v_i of data (e.g., few data points) can be highly valuable to the other parties if its data is distinct. Conversely, a party with a high value v_j of data should be less valuable to the other parties with similar data as its marginal contribution is lower.

Hence, we consider the Shapley value which captures the idea of marginal contribution precisely as it uses the expected marginal contribution of a party i when it joins the parties preceding it in any permutation:

$$\text{Shapley}_v(i) = \frac{1}{n!} \sum_{\pi \in \Pi_N} \left(v_{S_{\pi,i} \cup \{i\}} - v_{S_{\pi,i}} \right) \qquad (2)$$

[4] R2 and R3 are, respectively, adapted from $\forall C \in CS \;\; \sum_{i \in C} r_i \leq v_C$ and $\forall C \in CS \;\; \sum_{i \in C} r_i = v_C$ in CGT (Chalkiadakis et al., 2011).

[5] These properties are axioms of Shapley Value (Shapley, 1953) and have been widely used in existing ML works (Ghorbani & Zou, 2019; Jia et al., 2019b; Ohrimenko et al., 2019) for data valuation.

[6] Our definition is similar to coalitional monotonicity in (Young, 1985) except that we consider $>$ instead of $\geq$. This rules out the scenario where the value of party i's data improves but not that of its model reward. We further check the feasibility of improvement in the value of its model reward.

where Π_N is the set of all possible permutations of N and $S_{\pi,i}$ is the coalition of parties preceding i in permutation π.

Proposition 1. *$r_i = \text{Shapley}_v(i)$ for all $i \in N$ satisfy R5.*

Its proof is in Appendix B. A party will have a larger Shapley value and hence value r_i of model reward when its data is highly valuable on its own (e.g., with low inherent noise) and/or to the other parties (e.g., due to low correlation).

Fairness with Weak Efficiency (R3). We simplify the notation of $\text{Shapley}_v(i)$ to ϕ_i. Such a reward scheme may not satisfy R3 as the total value $\sum_{i\in N} r_i$ of model rewards is only v_N (Chalkiadakis et al., 2011). Due to the freely replicable nature of our model reward, we want the total value to exceed v_N and the value of some party's model reward to be v_N. To satisfy other incentive conditions, we consider a function g to map $(\phi_i)_{i\in N}$ to $(r_i)_{i\in N}$ (i.e., $r_i \triangleq g(\phi_i)$). To achieve fairness in R5, g must be strictly increasing with $g(0) = 0$. These motivate us to propose the following:

Definition 1 (Shapley Fairness). *Given $\{v_C\}_{C\in 2^N}$, if there exists a constant $k > 0$ s.t. $r_i = k\phi_i$ for all $i \in N$, then the values $(r_i)_{i\in N}$ of model rewards are Shapley fair.*

Note that CGT sets $k = 1$ (Proposition 1). Also, we can control constant k to satisfy other incentive conditions such as R3. A consequence of Definition 1 is that $\phi_i/\phi_j = r_i/r_j$. So, increasing the ratio of expected marginal contributions ϕ_i of party i vs. ϕ_j of party j (e.g., by including more valuable/informative data of party i, *ceteris paribus*) results in the same increase in the ratio of values of model rewards r_i received by party i vs. r_j received by party j.

Fairness with Individual Rationality (R4). However, another issue persists in the above modified definition of fairness (Definition 1): R4 may not be satisfied when the value of data is submodular (e.g., IG). We show an example below and leave the detailed discussion to Appendix C:

Example 1. *Suppose that there is a coalition of two parties whose values of data are $v_1 = 7$, $v_2 = 5$, and $v_{\{1,2\}} = 8$. Using (2), the Shapley values of the two parties are $\phi_1 = 5$ and $\phi_2 = 3$. To satisfy R3, we set $r_1 = 8$. Then, to achieve Shapley fairness, $r_2 = (8/5) \times 3 = 4.8$. Since $r_2 < v_2$, R4 is not satisfied.*

Since our model reward is freely replicable, it is possible to give all parties in a larger coalition more valuable model rewards. How then should we redefine g to derive larger values r_i of model rewards to achieve R4 while still satisfying R5 and R3 (i.e., $g(\max_{i\in N} \phi_i) = v_N$)? To answer this question, we further modify the above definition of Shapley fairness (Definition 1) to the following:

Definition 2 (ρ-Shapley Fairness). *Given $\{v_C\}_{C\in 2^N}$, if there exist constants $k > 0$ and $\rho > 0$ s.t. $r_i = k\phi_i^\rho$ for all $i \in N$, then the values $(r_i)_{i\in N}$ of model rewards are ρ-Shapley fair.*

It follows from Definition 2 that when $\rho < 1$, $\phi_i > \phi_j$ implies $\phi_i/\phi_j > r_i/r_j > 1$. Furthermore, reducing ρ from 1 "weakens" proportionality of r_i to ϕ_i and decreases the ratio of r_i vs. r_j while preserving $r_i > r_j$. So, if the value of a party's model reward is v_N (hence satisfying R3), then the values of the other parties' model rewards will become closer to v_N, thereby potentially satisfying R4. We will discuss the choice of k and the effect of ρ in Sec. 4.2.

4.1.2. Stability and Group Welfare

At the beginning of Sec. 4, we have provided justifications for desiring the grand coalition. We will now introduce two other solution concepts in CGT (i.e., stability and group welfare) to incentivize its formation and increase the total benefit from the collaboration, respectively. To the best of our knowledge, these solution concepts have not been considered in existing collaborative ML works.

A coalition structure CS with a given set $(r_i)_{i\in N}$ of values of model rewards is said to be *stable* if no subset of parties has a common incentive to abandon it to form another coalition on their own. With stability, parties can be assured that the collaboration will not fall apart and they cannot get a more valuable model reward from adding or removing parties. Similar to other solution concepts in CGT introduced previously, the definition of stability in CGT[7] does not suit the freely replicable nature of our model reward and the redefined incentive condition on feasibility R2. Therefore, we propose the following new definition of stability:

Definition 3 (Stability). *A coalition structure CS with a given set $(r_i)_{i\in N}$ of values of model rewards is stable if $\forall C \subseteq N \;\; \exists i \in C \;\; v_C \leq r_i$.*

Conversely, supposing $\exists C \subseteq N \;\; \forall i \in C \;\; r_i < v_C$, all parties in C may be willing to deviate to form coalition C as they can feasibly increase the values of their model rewards (up) to v_C. The condition in Definition 3 is computationally costly to check as it involves an exponential (in n) number of constraints. To ease computation and since we are mainly interested in the grand coalition (i.e., $CS = \{N\}$), we look at the following sufficient condition instead:

R6 **Stability of Grand Coalition.** Suppose that the value of data is monotonic. The grand coalition is stable if for every coalition C, the value of the model reward received by the party with largest Shapley value is at least v_C:
$\forall C \subseteq N \;\; \forall i \in C \;\; \phi_i = \max_{j\in C} \phi_j \Rightarrow \; v_C \leq r_i$.

The values $(r_i)_{i\in N}$ of model rewards that satisfy R6 will also satisfy Definition 3. To ensure R6, for each party $i \in N$, we will only need to check one constraint involving r_i:

[7] In CGT, a coalition structure CS with given $(r_i)_{i\in N}$ is *stable* if $\forall C \subseteq N \;\; v_C \leq \sum_{i\in C} r_i$.

$v_{C_i} \leq r_i$ where C_i includes all parties whose Shapley value is at most ϕ_i. The monotonic property of the value of data guarantees that $v_C \leq v_{C_i} \leq r_i$ for any $C \subseteq C_i$. The total number of constraints is linear in the number n of parties. Unlike R1 to R5, R6 is an optional incentive condition.

The last incentive condition stated below does not have to be *fully* provided:

R7 **Group Welfare.** The values $(r_i)_{i \in N}$ of model rewards should maximize the group welfare $\sum_{i \in N} r_i$.

4.2. Reward Scheme Considering All Incentives

In this subsection, we will present a reward scheme which considers all the incentives in Sec. 4.1 and assumes that the grand coalition will form. It uses parameter ρ in Definition 2 to trade off between achieving Shapley fairness vs. satisfying the other incentives, as detailed below:

Theorem 1. *Let $0 < \rho \leq 1$. For each party $i \in N$, let $\phi_i \triangleq \text{Shapley}_v(i)$ and reward $r_i \triangleq (\phi_i/\phi^*)^\rho \times v_N$ where $\phi^* = \max_{i \in N} \phi_i$.*[8] *The values $(r_i)_{i \in N}$ of model rewards are ρ-Shapley fair and satisfy R1 to R3 and R5 when $\rho > 0$. Also, when*

- *$\rho = 1$, $(r_i)_{i \in N}$ are (pure) Shapley fair (Definition 1);*
- *$\rho \leq \rho_r \triangleq \min_{i \in N} \log(v_i/v_N)/\log(\phi_i/\phi^*)$, $(r_i)_{i \in N}$ satisfy individual rationality (R4);*
- *$\rho \leq \rho_s \triangleq \min_{i \in N} \log(v_{C_i}/v_N)/\log(\phi_i/\phi^*)$ where coalition $C_i \triangleq \{j \in N \mid \phi_j \leq \phi_i\}$, $(r_i)_{i \in N}$ achieve stability of the grand coalition (R6) and individual rationality (R4) as $\rho_s \leq \rho_r$;*
- *$\rho = 0$, $(r_i)_{i \in N}$ provide maximum group welfare (R7) but do not satisfy fairness (R5).*

Its proof is in Appendix D. As ρ decreases from 1, the values $(r_i)_{i \in N}$ of model rewards deviate further from pure Shapley fairness (i.e., less proportional to $(\phi_i)_{i \in N}$) and the value r_i of any party i's model reward with $0 < \phi_i/\phi^* < 1$ will increase. This increases group welfare (R7) and the values $(r_i)_{i \in N}$ of model rewards can potentially satisfy individual rationality (R4) and stability (R6). Thus, a smaller ρ addresses the limitations of pure Shapley fairness. We do not consider (a) $\rho > 1$ as it reduces group welfare and is not needed to satisfy other incentives, nor (b) $\rho < 0$ as it demands that more valuable/informative data and higher Shapley value lead to less valuable model reward, hence not satisfying fairness (R5).

Fig. 1 gives an overview of the incentive conditions and how Theorem 1 uses ρ to satisfy them and trade off between the various incentives. Note that R6 guarantees R4. Ideally, we want the values $(r_i)_{i \in N}$ of model rewards to satisfy fairness

[8] In practice, ϕ_i can be based on other solution concepts in CGT that satisfy F1 to F4.

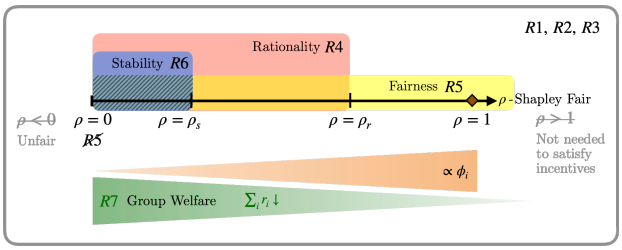

Figure 1. An overview of the incentive conditions and how they are satisfied by our reward scheme in Theorem 1. There may be alternative desirable reward schemes not using Definition 2. Note that in some scenarios, 1 may be less than ρ_r or ρ_s. As a result, the values $(r_i)_{i \in N}$ of model rewards that are purely Shapley fair naturally satisfy individual rationality and stability, respectively.

(R5), stability (R6), and individual rationality (R4) (i.e., shaded region in Fig. 1). However, the values $(r_i)_{i \in N}$ of model rewards that are purely Shapley fair may not always lie in the desired shaded region. So, a smaller ρ is needed.

Consider how much each party i benefits from the collaboration by choosing a ρ smaller than 1, specifically, by measuring the ratio of the values r_i of its received model reward that are ρ-Shapley fair ($\rho < 1$) vs. purely Shapley fair ($\rho = 1$). Such a ratio can be evaluated to $(\phi^*/\phi_i)^{1-\rho} \geq 1$ and is smaller (larger) for a party i with a larger (smaller) ϕ_i. So, when a smaller ρ is chosen, a party i with ϕ_i close to ϕ^* needs to be "altruistic" as it cannot benefit as much due to a smaller ratio than any party j with a smaller ϕ_j. However, note that party i is already getting close to the maximum value v_N of model reward. Regardless of the chosen ρ, the party i with the largest ϕ_i (i.e., ϕ^*) always receives the most valuable model reward with the highest value v_N. In practice, the parties may agree to use a smaller ρ if they want to increase their total benefit from the collaboration or if they do not know their relative expected marginal contributions beforehand.

4.3. Realization of Model Rewards

Finally, we will have to train a model for each party as a reward to realize the values $(r_i^*)_{i \in N}$ of model rewards decided by the above reward scheme. Recall from the beginning of Sec. 4 that the value $r_i = \mathbb{I}(\boldsymbol{\theta}; D_i^r)$ of model reward received by each party i is the IG on model parameters θ given the data D_i^r. How is the training data D_i^r for each party i selected to realize $r_i = r_i^*$?

A direct approach is to select and only train on a subset of the aggregated data from all parties: $D_i^r \subseteq D_N$. However, this is infeasible due to the need to consider an *exponential* number of *discrete* subsets of data, all of which may not be able to realize the decided value r_i^* of model reward exactly.

To avoid the above issue, we instead consider injecting Gaussian noise $\mathcal{N}(\mathbf{0}, \eta_i \mathbf{I})$ into the aggregated data from multiple parties and optimizing the *continuous* noise variance param-

eter η_i for realizing the decided value r_i^* of model reward. Specifically, the central party collects data $D_i \triangleq (\mathbf{X}_i, \mathbf{y}_i)$ from every party $i \in N$ and constructs D_i^r by concatenating the input matrices $\mathbf{X}_i$ for $i \in N$ and the output vector $\mathbf{y}_i$ with $\mathbf{z}_i \triangleq (\mathbf{y}_j)_{j \in N \setminus \{i\}} + \mathcal{N}(\mathbf{0}, \eta_i \mathbf{I})$. When training a model for party i, we use the original $\mathbf{y}_i$ so that each party gets all the information from its own data D_i and cannot improve its received model reward by subsequently training on it. We inject noise with variance η_i only to the other parties' data (i.e., $(\mathbf{y}_j)_{j \in N \setminus \{i\}}$) to reduce the information from parties $N \setminus \{i\}$ conveyed to party i. In particular, $r_i = v_N$ when $\eta_i = 0$ and $r_i = v_i$ when $\eta_i = \infty$. By varying η_i, we can span different values of r_i. We use an efficient root-finding algorithm to find the optimal η_i such that $r_i = r_i^*$ for each party $i \in N$. The effect of adding noise to the data will be reflected in the variance of party i's model reward parameters and its predictive distribution. We will show that the added noise affects predictive accuracy reasonably in Sec. 5.

5. Experiments and Discussion

This section empirically evaluates the performance and properties of our reward scheme (Sec. 4.2) on Bayesian regression models with the (a) synthetic Friedman dataset with 6 input features (Friedman, 1991), (b) *diabetes progression* (DiaP) dataset on the diabetes progression of 442 patients with 9 input features (Efron et al., 2004), and (c) *Californian housing* (CaliH) dataset on the value of 20640 houses with 8 input features (Pace & Barry, 1997). We use a Gaussian likelihood and assume that the model hyperparameters are known or learned using maximum likelihood estimation.

The performance of our reward scheme is evaluated using the IG and *mean negative log probability* (MNLP) metrics:

$$\begin{aligned}\text{MNLP} &\triangleq \frac{1}{|D^*|} \sum_{(\mathbf{x}^*, y^*) \in D^*} -\log p(y^*|\mathbf{x}^*, D_i^r) \\ &= \frac{1}{|D^*|} \sum_{(\mathbf{x}^*, y^*) \in D^*} \frac{1}{2}\left(\log(2\pi\sigma_*^2) + \frac{(\mu_* - y^*)^2}{\sigma_*^2}\right)\end{aligned} \tag{3}$$

where D^* denotes a test dataset, and μ_* and σ_*^2 denote, respectively, the predictive mean and variance of the predictive distribution $p(y^*|\mathbf{x}^*, D_i^r)$. We then use these metrics to show how an improvement in IG can affect the predictive accuracy of a trained model on a common test dataset.[9] More details of the experimental settings such as how to select the test dataset are in Appendix E. The performance of our reward scheme cannot be directly compared against that of the existing works (Ghorbani & Zou, 2019; Jia et al., 2019b) as they mainly focus on non-Bayesian classification models and assume monetary compensations. For all experiments, we consider a collaboration of $n = 3$ parties as it suffices to show interesting results.

[9] We use a randomly sampled test dataset to illustrate how IG is a suitable surrogate measure of the test/predictive accuracy of a trained model. In Sec. 3, we have discussed that in practice, it is often tedious or even impossible for all parties to agree on a common validation dataset, let alone a common set of test queries.

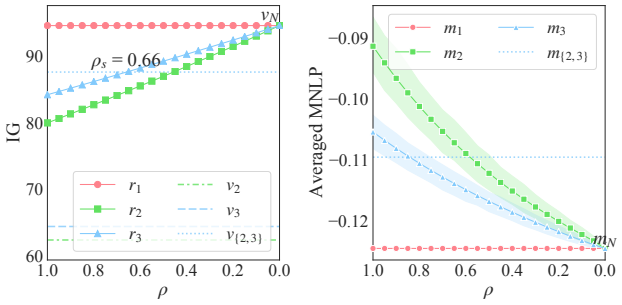

(a) IG (b) MNLP

Figure 2. Graphs of (a) IG and (b) MNLP vs. the adjustable parameter ρ for GP regression with synthetic Friedman dataset where m_i and m_C (with $|C| > 1$) denote the MNLPs of the models, respectively, rewarded to party i and trained using data D_C.

Gaussian process (GP) regression with synthetic Friedman dataset. For each party, we generate 250 data points from the Friedman function. We assign party 1 the most valuable data by spanning the first input feature over the entire input domain $[0, 1]$. In contrast, for parties 2 and 3, the same input feature spans only the non-overlapping ranges of $[0, 0.5]$ and $[0.5, 1]$, respectively. This makes the data of parties 2 and 3 highly valuable to each other, i.e., $v_{\{2,3\}}$ is high relative to v_2. Using (2), the Shapley values of the three parties are $\phi_1 = 34.57$, $\phi_2 = 29.24$, and $\phi_3 = 30.78$. Party 1 has the largest ϕ_1 and will always receive the most valuable model reward with the highest IG or value v_N.

Fig. 2a shows that our reward scheme indeed satisfies fairness (R5): A party i with a larger ϕ_i always receives a higher r_i (i.e., more valuable model reward). Also, since $\rho_r \geq 1$ here, our reward scheme always satisfies individual rationality (R4): Every party i receives a more valuable model reward than that trained on its own data, i.e., its IG r_i is higher than v_i (see dotted lines in Fig. 2a). As ρ decreases (rightwards), party 1's most valuable model reward is unaffected but it has to be "altruistic" to parties 2 and 3 with increasing r_2 and r_3 (i.e., increasingly valuable model rewards) but smaller ϕ_2 and ϕ_3, as discussed in the last paragraph of Sec. 4.2. When ρ decreases to ρ_s, stability (R6) is reached: Party 3's model reward matches what it will receive by only collaborating with party 2. When $\rho = 0$, all parties receive an equally valuable model reward with the same IG/value v_N despite their varying expected marginal contributions. This shows how ρ can be reduced to trade off proportionality in pure Shapley fairness for satisfying other incentives such as stability (R6) and group welfare (R7).

In Fig. 2b, we report the averaged MNLP and shade the 95% confidence interval over 20 different realizations of $\mathbf{z}_i$. As ρ

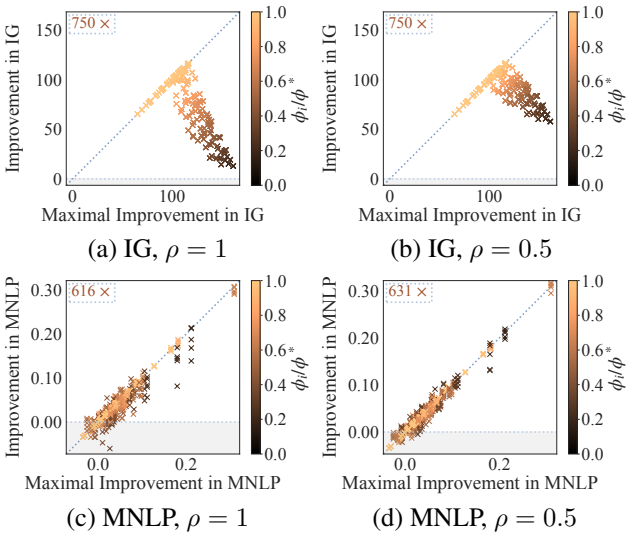

(a) IG, $\rho = 1$ (b) IG, $\rho = 0.5$

(c) MNLP, $\rho = 1$ (d) MNLP, $\rho = 0.5$

Figure 3. Scatter graph of the (a-b) improvement in IG ($r_i - v_i$) vs. maximal improvement in IG ($v_N - v_i$) and (c-d) improvement in MNLP vs. maximal improvement in MNLP when $\rho = 1, 0.5$ for multiple partitions of DiaP dataset among $n = 3$ parties.

decreases and each party receives an increasingly valuable model reward (i.e., increasing IG/r_i), MNLP decreases, thus showing that an improvement in IG translates to a higher predictive accuracy of its trained model.

GP regression with DiaP dataset. We train a GP regression model with a composite kernel (i.e., squared exponential kernel + exponential kernel) and partition the training dataset among $n = 3$ parties, as detailed later. The results are shown in Fig. 3.

Neural network regression with CaliH dataset. We consider *Bayesian linear regression* (BLR) on the last layer of a *neural network* (NN). 60% of the CaliH data is designated as the "public" dataset[10] and used to pretrain a NN. Since the correlation of the house value with the input features in the data of the parties may differ from that in the "public" dataset, we perform transfer learning and selectively retrain only the last layer using BLR with a standard Gaussian prior. So, IG is only computed based on BLR. From the remaining data, we select 400 data points for the test dataset and 1600 data points[11] for the training dataset to be partitioned among 3 parties, as detailed later. The results are shown in Fig. 5.

Sparse GP regression with synthetic Friedman dataset (10 parties). We also evaluate the performance of our reward scheme on the collaboration between a larger number $n = 10$ of parties. We consider a larger synthetic Friedman dataset of size 5000 and partition the training dataset among 10 parties such that each party owns at least 5% of the dataset. We train a sparse GP model based on the *deterministic training conditional* approximation (Hoang et al., 2015; 2016) and a squared exponential kernel. Since the exact Shapley value (2) is expensive to evaluate with a large number of parties, we approximate it using the *simple random sampling* algorithm in (Maleki et al., 2013). The results are shown in Fig. 6.

[10] In practice, the "public" dataset can be provided by the government to kickstart the collaborative ML. Alternatively, it can be historical data that companies are more willing to share.

[11] We restrict the total number of data points so that every individual party cannot train a high-quality model alone.

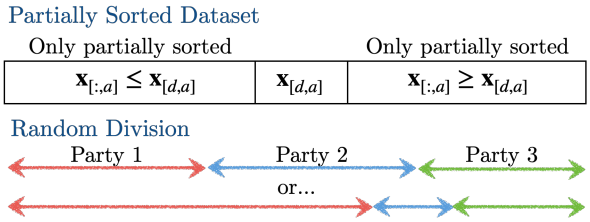

Figure 4. Only data point d will be at its sorted position based on input feature a. We then divide this partially sorted training dataset into 3 consecutive blocks of random sizes with the constraint that each party owns at least 10% of the dataset.

Partitioning Algorithm. To divide any training dataset among parties and vary their contributed data, we first choose an input feature a uniformly at random from all the input features. Next, the data point d is chosen uniformly at random. We partially sort the dataset and divide it into continuous blocks of random sizes, as detailed in Fig. 4. Such a setup allows parties to have different quantities of unique or partially overlapping data. The input feature a may have real-world significance: If a is the age of patients, then this models the scenario where hospitals have data of patients with different demographics.

Summary of Observations. For any dataset, we consider different train-test splits and partitions of the training dataset using the above algorithm. For each partition, we compute the optimized noise variance η_i to realize r_i^* and consider 5 realizations of $\mathbf{z}_i$ for each party $i \in N$. For a given partition and choice of ρ, if the values $(r_i)_{i \in N}$ of model rewards satisfy individual rationality (R4), then we compute the ratio ϕ_i/ϕ^*, improvement in IG ($r_i - v_i$) and MNLP, and maximal (possible) improvement in IG ($v_N - v_i$) and MNLP for each party $i \in N$. The improvement is measured relative to a model trained only on party i's data. The best/lowest MNLP m_N is assumed to be achieved by a model trained on the aggregated data of all parties. Each realization of $\mathbf{z}_i$ (from each party i) will produce a point in the scatter graphs in Figs. 3, 5, and 6. If a point lies on the diagonal identity line, then the corresponding party (e.g., with the largest ϕ_i, i.e., ϕ^*) receives a model reward with MNLP m_N. We also report the number of points with positive improvement in IG and MNLP in the top left hand corner of each graph.

In Figs. 3c-d, 5c-d, and 6c-d, the improvement in MNLP is usually positive: The predictive performance of each party benefits from the collaboration. Occasionally but reasonably,

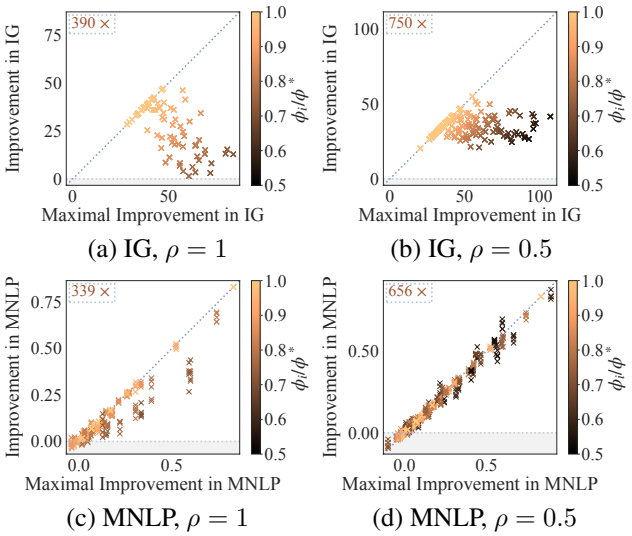

(a) IG, $\rho = 1$ (b) IG, $\rho = 0.5$

(c) MNLP, $\rho = 1$ (d) MNLP, $\rho = 0.5$

Figure 5. Scatter graph of the (a-b) improvement in IG ($r_i - v_i$) vs. maximal improvement in IG ($v_N - v_i$) and (c-d) improvement in MNLP vs. maximal improvement in MNLP when $\rho = 1, 0.5$ for multiple partitions of CaliH dataset among $n = 3$ parties.

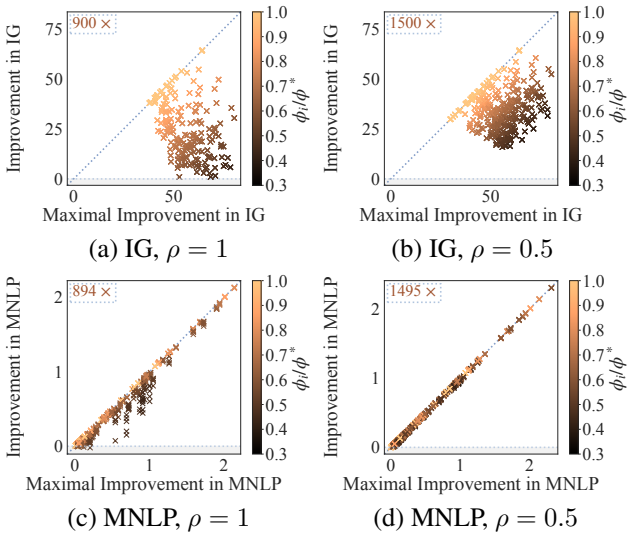

(a) IG, $\rho = 1$ (b) IG, $\rho = 0.5$

(c) MNLP, $\rho = 1$ (d) MNLP, $\rho = 0.5$

Figure 6. Scatter graph of the (a-b) improvement in IG ($r_i - v_i$) vs. maximal improvement in IG ($v_N - v_i$) and (c-d) improvement in MNLP vs. maximal improvement in MNLP when $\rho = 1, 0.5$ for multiple partitions of Friedman dataset among $n = 10$ parties.

this does not hold when the maximal improvement in MNLP is extremely small or even negative. Lighter colored points lie closer to the diagonal line: Parties with larger ϕ_i receive more valuable model rewards (Figs. 3a-b, 5a-b, 6a-b) which translates to lower MNLP (Figs. 3c-d, 5c-d, 6c-d).

In Figs. 3, 5 and 6, when ρ decreases from 1 to 0.5, darker colored points move closer to the diagonal line, which implies that parties with smaller ϕ_i can now receive more valuable model rewards with higher predictive accuracy. The number of points (reported in the top left corner) also increase as more of them satisfy R4.

In Fig. 6d with $\rho = 0.5$, most points are close to the diagonal identity line because it may be the case that not all training data points are needed to achieve the lowest MNLP. Instead, fewer or noisier data points are sufficient. In this scenario, a larger ρ may be preferred to reward the parties differently. More experimental results for different Bayesian models and datasets are shown in Appendix F.1.

Limitations. Though IG is a suitable surrogate measure of the predictive accuracy of a trained model, a higher IG does not always translate to lower MNLP. In Figs. 3c-d and 5c-d, occasionally, the improvement in MNLP is negative (around 15% of the time) and darker points lie above the diagonal line. The latter suggests that parties with smaller ϕ_i may receive model rewards with MNLP lower than m_N (i.e., awarded to the party with largest ϕ_i, i.e., ϕ^*). This may be purely due to randomness, but we have also investigated factors (e.g., model selection) leading to a weak relationship between IG and MNLP and reported the results in Appendix F.2. Our reward scheme works best when a suitable model is selected and the model prior is not sufficiently informative for any party to achieve a high predictive accuracy by training a model on its own data. In practice, this is reasonable as collaboration precisely happens only when every individual party cannot train a high-quality model alone but can do so from working together.

6. Conclusion

This paper describes a collaborative ML scheme that distributes only model rewards to the parties. We perform data valuation based on the IG on the model parameters given the data and compute each party's marginal contribution using the Shapley value. To decide the appropriate rewards to incentivize the collaboration, we adapt solution concepts from CGT (i.e., fairness, Shapley fairness, stability, individual rationality, and group welfare) for our freely replicable model rewards. We propose a novel reward scheme (Theorem 1) with an adjustable parameter that can trade off between these incentives while maintaining fairness. Empirical evaluations show that a smaller ρ and more valuable model rewards translate to higher predictive accuracy. Current efforts on designing incentive mechanisms for federated learning can build upon our work presented in this paper. For future work, we plan to address privacy preservation in our reward scheme. The noise injection method used for realizing the model rewards in this work is closely related to the Gaussian mechanism of differential privacy (Dwork & Roth, 2014). This motivates us to explore how the injected noise will affect privacy in the model rewards.

Acknowledgements

This research/project is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative and its AI Singapore Programme (Award Number: AISG-GC-2019-002). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

Center for Open Data Enterprise. Sharing and utilizing health data for AI applications. Roundtable report, 2019.

Chalkiadakis, G., Elkind, E., and Wooldridge, M. Computational aspects of cooperative game theory. In Brachman, R. J., Cohen, W. W., and Dietterich, T. G. (eds.), *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2011.

Claver, H. Data sharing key for AI in agriculture. Future Farming, Feb. 2019. URL https://www.futurefarming.com/Tools-data/Articles/2019/2/Data-sharing-key-for-AI-in-agriculture-389844E/.

Conway, J. *Artificial Intelligence and Machine Learning: Current Applications in Real Estate*. PhD thesis, Massachusetts Institute of Technology, 2018.

Cover, T. and Thomas, J. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.

de Clippel, G. and Rozen, K. Fairness through the lens of cooperative game theory: An experimental approach. Cowles Foundation discussion paper no. 1925, 2013.

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004.

Friedman, J. H. Multivariate adaptive regression splines. *Ann. Statist.*, 19(1):1–67, 1991.

Ghorbani, A. and Zou, J. Data Shapley: Equitable valuation of data for machine learning. In *Proc. ICML*, pp. 2242–2251, 2019.

Ghorbani, A., Kim, M. P., and Zou, J. A distributional framework for data valuation. In *Proc. ICML*, 2020.

Harrison, D. and Rubinfeld, D. Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management*, 5(1):81–102, 1978.

Hoang, T. N., Hoang, Q. M., and Low, K. H. A unifying framework of anytime sparse Gaussian process regression models with stochastic variational inference for big data. In *Proc. ICML*, pp. 569–578, 2015.

Hoang, T. N., Hoang, Q. M., and Low, K. H. A distributed variational inference framework for unifying parallel sparse Gaussian process regression models. In *Proc. ICML*, pp. 382–391, 2016.

Jia, R., Dao, D., Wang, B., Hubis, F. A., Gurel, N. M., Li, B., Zhang, C., Spanos, C., and Song, D. Efficient task-specific data valuation for nearest neighbor algorithms. *Proc. VLDB Endowment*, 12(11):1610–1623, 2019a.

Jia, R., Dao, D., Wang, B., Hubis, F. A., Hynes, N., Gurel, N. M., Li, B., Zhang, C., Song, D., and Spanos, C. Towards efficient data valuation based on the Shapley value. In *Proc. AISTATS*, pp. 1167–1176, 2019b.

Krause, A. and Guestrin, C. Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach. In *Proc. ICML*, pp. 449–456, 2007.

Kruschke, J. K. Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior*, 36(3):210–226, 2008.

Maleki, S., Tran-Thanh, L., Hines, G., Rahwan, T., and Rogers, A. Bounding the estimation error of sampling-based Shapley value approximation. arXiv:1306.4265, 2013.

Maschler, M. and Peleg, B. A characterization, existence proof and dimension bounds for the kernel of a game. *Pacific J. Mathematics*, 18(2):289–328, 1966.

Ocean Protocol Foundation. Ocean protocol: A decentralized substrate for AI data and services. Technical whitepaper, 2019.

Ohrimenko, O., Tople, S., and Tschiatschek, S. Collaborative machine learning markets with data-replication-robust payments. arXiv:1911.09052, 2019.

Pace, R. K. and Barry, R. Sparse spatial auto-regressions. *Statistics and Probability Letters*, 33(3):291–297, 1997.

Shapley, L. S. A value for n-person games. In Kuhn, H. W. and Tucker, A. W. (eds.), *Contributions to the Theory of Games*, volume 2, pp. 307–317. Princeton Univ. Press, 1953.

Yoon, J., Arik, S. O., and Pfister, T. Data valuation using reinforcement learning. In *Proc. ICML*, 2020.

Young, H. P. Monotonic solutions of cooperative games. *International Journal of Game Theory*, 14(2):65–72, 1985.