

DUPRE: Data Utility Prediction for Efficient Data Valuation

Kieu Thao Nguyen Pham*
National University of Singapore
Singapore
nguyen.pkt@u.nus.edu

Rachael Hwee Ling Sim*
National University of Singapore
Singapore
rachael.sim@u.nus.edu

Quoc Phong Nguyen
A2I2, Deakin University
Australia
qphongmp@gmail.com

See Kiong Ng
National University of Singapore
Singapore
seekiong@nus.edu.sg

Bryan Kian Hsiang Low
National University of Singapore
Singapore
lowkh@comp.nus.edu.sg

ABSTRACT

Data valuation is increasingly used in machine learning (ML) to decide the fair compensation for data owners and identify valuable or harmful data for improving ML models. Cooperative game theory-based data valuation, such as Data Shapley, requires evaluating the data utility (e.g., validation accuracy) and retraining the ML model for multiple data subsets. While most existing works on efficient estimation of the Shapley values have focused on reducing the number of subsets to evaluate, our framework, DUPRE, takes an alternative yet complementary approach that reduces the cost per subset evaluation by predicting data utilities instead of evaluating them by model retraining. Specifically, given the evaluated data utilities of some data subsets, DUPRE fits a *Gaussian process* (GP) regression model to predict the utility of every other data subset. Our key contribution lies in the design of our GP kernel based on the sliced Wasserstein distance between empirical data distributions. In particular, we show that the kernel is valid and positive semi-definite, encodes prior knowledge of similarities between different data subsets, and can be efficiently computed. We empirically verify that DUPRE introduces low prediction error and speeds up data valuation for various ML models, datasets, and utility functions.

KEYWORDS

Data Valuation; Gaussian Process Regression; Shapley Value; Semivalue; Data Utility Prediction; Kernel; Collaborative Machine Learning

ACM Reference Format:

Kieu Thao Nguyen Pham, Rachael Hwee Ling Sim, Quoc Phong Nguyen, See Kiong Ng, and Bryan Kian Hsiang Low. 2025. DUPRE: Data Utility Prediction for Efficient Data Valuation. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

*Equal contribution.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

In recent years, there has been growing interest in *data valuation* and understanding how much data is worth in *machine learning* (ML). Data valuation can be used to determine the fair compensation that data owners deserve for sharing their data [9, 27] and to identify valuable datasets to explain and improve the performance of their models [8]. A common category of data valuation methods that values a data point/set (relative to the data contributed by others) is *cooperative game theory* (CGT) based valuation [26]. Suppose the ML model is trained on data from a set N of n data owners (owners). Data Shapley [8], a popular CGT-based valuation technique, values an owner i by its Shapley value

$$\phi_u(i) \triangleq \sum_{C \subseteq N \setminus i} \frac{1}{n} \binom{n-1}{|C|} [u(C \cup \{i\}) - u(C)]. \quad (1)$$

The data utility function u maps any coalition (i.e., set) $C \subseteq N$ of owners to the ML performance achievable by their data. A concrete example of u is the validation accuracy on a trained neural network. Other CGT-based valuations include the least-core solution [34] and semivalues [18, 30], such as the Banzhaf value, which is similar to the Shapley value but uses a different set of weights $(\omega_c)_{c=0}^{n-1}$ such that each $\omega_c \geq 0$ and $\sum_{c=0}^{n-1} \omega_c \binom{n-1}{c} = 1$ in the following expression

$$\phi_u(i) \triangleq \sum_{C \subseteq N \setminus i} \omega_{|C|} [u(C \cup \{i\}) - u(C)]. \quad (2)$$

While CGT-based valuations satisfy desirable axioms, they all require an exponential number of evaluations of u . This high complexity motivates existing works to study more efficient methods for estimating the Shapley value and other semivalues by (i) reducing the number of evaluations or (ii) reducing the cost per evaluation. Each data utility evaluation may involve expensive model retraining from scratch to evaluate its predictive performance. Our work addresses (ii) and is complementary to methods that address (i).

Most existing works focus on (i) and propose Monte Carlo methods, such as permutation sampling, group testing [9], and reusing samples (data utility evaluations) to compute the semivalues for multiple owners efficiently [16, 19]. To address (ii), some works propose heuristics to avoid retraining the ML model from scratch. For example, TMC-Shapley [8] simply approximates $\Delta_u(i|C) \triangleq u(C \cup i) - u(C)$ with 0 when $u(C)$ is sufficiently close to $u(N)$. Gradient Shapley [8] considers training the ML model (e.g., deep neural network) over a single training epoch so multiple utilities

(e.g., $u(C)$, $u(C \cup i)$ then $u(C \cup \{i, j\})$) can be incrementally computed. Jia et al. [9] suggest using the influence function heuristic to approximate $\Delta_u(i|C)$ when owner i owns a single data point. It also describes how to compute the Shapley value of all points exactly in log-linear time for their k -nearest neighbor utility function. However, these methods cannot be applied for *all* models (e.g., neural networks trained over *multiple* epochs), utility functions (e.g., when $u(C) \ll u(N)$) and datasets (e.g., the influence approximation may be inaccurate when each owner owns a larger dataset instead [15]). These limitations raise an important question: For faster data valuation, is there a *general* method to *reduce the cost per data utility evaluations* that works for *all* models, utility functions, and datasets?

Wang et al. [31] have proposed predicting data utility for any input dataset and using a hybrid of actual utility evaluations and predicted evaluations during data valuation. They briefly analyzed how this hybrid slightly worsens the CGT-based valuation approximation error. However, Wang et al. [31] did not optimize the predictor and simply trained a neural network that takes in an indicator vector for each coalition C (i.e., an n -dimensional vector where each entry is 1 if the corresponding owner is present in C and 0 if absent).

Our work seeks to optimize the predictor and addresses the following questions: (1) Can the predictor further exploit the similarity between data of different owners? For example, if owners i and j have highly similar data, can the predictor leverage the prior knowledge that $u(C \cup i) \approx u(C \cup j)$ (instead of learning the relationship from more data utility evaluations)? (2) Additionally, can the predictor quantify the uncertainty in its prediction?

Gaussian process regression (GPR) seems well-suited to this task. We can (1) incorporate prior knowledge by specifying an appropriate kernel over datasets and (2) quantify the additional uncertainty in the estimated semivalue (due to the predictor) using the GPR posterior covariance. However, adapting GPR presents challenges: The kernel must be *positive semi-definite* (PSD), to ensure a nondegenerate GP posterior and offer computational savings as compared to directly evaluating the utility from model retraining. We propose measuring the similarity between empirical data distributions with a *sliced Wasserstein distance* (SW) kernel, as Meunier et al. [22] have proven that the kernel is PSD. Additionally, our method is computationally efficient: after precomputing the sorted projections once, the SW distance used in each kernel entry can be computed in linear time w.r.t. the total dataset size. Our method, DUPRE, fits a GP model based on some actual data utility evaluations. Subsequently, DUPRE predicts the utility of unevaluated coalitions and can be used to identify those with high uncertainty for further actual evaluations. DUPRE complements any exact CGT-based valuation and approximation techniques proposed to address (i).

In summary, we make the following key contributions. In Section 2, we formulate the problem of predicting the utility of coalitions using GPR and propose a suitable valid kernel to measure the similarity between datasets for regression and classification problems. In Section 3, we describe how to estimate the semivalues based on the GP model.¹ In Section 4, we empirically verify

that DUPRE introduces a low prediction error and speeds up data valuation for various few models, datasets, and utility functions.

2 PROBLEM FORMULATION

We consider n data owners $N = \{1, \dots, n\}$. Data owner i has a dataset $D_i \triangleq (X_i, y_i)$ where X_i is the input matrix and y_i is the target outputs. Each owner i shares their dataset D_i with the mediator, who train ML models on data from multiple owners to assign a fair data value ϕ_i to each data owner i . Let $C \subseteq N$ denote a *coalition* of data owners with the aggregated dataset $D_C \triangleq (X_C, y_C)$, where $X_C \triangleq \cup_{i \in C} X_i$ and $y_C \triangleq \cup_{i \in C} y_i$.²

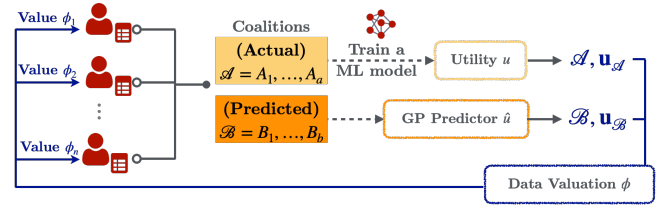


Figure 1: We partition all (sampled) coalitions into two families of sets. The utilities of coalitions in \mathcal{A} are actually evaluated by training the ML model and measuring the utility, e.g., validation accuracy. In contrast, the utilities of coalitions in \mathcal{B} are predicted by the Gaussian Process (GP) model.

The data utility function u maps any coalition C to the performance (e.g., negated mean squared error) of the ML model trained on their data, $u(C)$. The function u may be expensive to evaluate for complex models such as deep neural networks.

For any j , let coalitions A_j and B_j be subsets of N . Given a actual utility evaluations of coalitions in $\mathcal{A} \triangleq (A_j)_{j=1}^a$ (i.e., $(u(A))_{A \in \mathcal{A}}$), our goal is to learn a predictor \hat{u} that predicts b data utilities of coalitions in $\mathcal{B} \triangleq (B_j)_{j=1}^b$. Subsequently, we use both the actual utility evaluations at \mathcal{A} and predicted utility evaluations at \mathcal{B} to exactly compute or approximate the CGT-based data valuation, as shown in Figure 1.

For each coalition C , we model the data utility $u(C) = v(C) + \epsilon_C$ where ϵ_C is sampled from a Gaussian distribution with zero mean and variance σ^2 . We specify the underlying generating function v as a *Gaussian process* [32] with a covariance kernel k (see Appendix A.1). Given the observed utilities $\mathbf{u}_{\mathcal{A}}$ for coalitions \mathcal{A} , the posterior belief of the utilities $\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}$ for coalitions \mathcal{B} follows a Gaussian distribution

$$\begin{aligned} \hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}} &\sim \mathcal{N}\left(\mathbb{E}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}], \mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}]\right) \\ \mathbb{E}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}] &= \mathbf{K}_{\mathcal{B},\mathcal{A}}[\mathbf{K}_{\mathcal{A},\mathcal{A}} + \sigma^2\mathbf{I}]^{-1}\mathbf{u}_{\mathcal{A}} \\ \mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}] &= \mathbf{K}_{\mathcal{B},\mathcal{B}} - \mathbf{K}_{\mathcal{B},\mathcal{A}}[\mathbf{K}_{\mathcal{A},\mathcal{A}} + \sigma^2\mathbf{I}]^{-1}\mathbf{K}_{\mathcal{A},\mathcal{B}}. \end{aligned}$$

Here, $\mathbf{K}_{\mathcal{B},\mathcal{A}}$ is a matrix whose j, i entry is $k(B_j, A_i)$, the similarity between the aggregated datasets of D_{B_j} and that of D_{A_i} . Notice that there are $O(a^2 + ab + b^2)$ unique kernel entries in $\mathbf{K}_{\mathcal{B},\mathcal{A}}$ and $\mathbf{K}_{\mathcal{A},\mathcal{A}}$. The inverse of the kernel matrix only needs to be computed once

¹The supplementary materials can be found at <https://kakaeriol.github.io/dupre/>.

²For notation convenience, we say $(x, y) \in D_C$ if there exists an index j such that the j element of X_C and y_C are x and y , respectively.

in $O(a^3)$ time. Subsequently, each utility prediction only involves matrix multiplication.

Our next challenge is to decide the kernel function k over datasets or data distributions such that the kernel is **(I)** valid (see Appendix B for properties a valid kernel must satisfy such as positive semi-definite, PSD) and **(II)** computationally efficient. The former results in a valid GP while the latter ensures that the method is useful in practice.

2.1 Choice of Kernel

For simplicity, we start by ignoring the target outputs and considering only the input matrix, i.e., $D_i = X_i$ for every i . How do we measure the similarity between the aggregated dataset D_A (from owners in A) and D_B ? Equivalently, let δ be the Dirac delta distribution, how do we measure the distance between the empirical data distributions $\mathbb{P}(D_A) = \frac{1}{|D_A|} \sum_{x \in D_A} \delta(x)$ and $\mathbb{P}(D_B)$? We quantify the distance between the empirical data distributions using optimal transport distances (see Appendix A.2) rather than f-divergences. Optimal transport distances, such as the Wasserstein distance, exhibit desirable mathematical properties including symmetry and the triangle inequality, which are essential for defining valid kernels and comparing distributions even when their supports are disjoint. Optimal transport distances measure the minimal total cost required to transform one distribution into another. In particular, the Wasserstein distance captures the intrinsic geometry of the space of distributions [22] and admits an intuitive interpretation: it is the minimum total cost of transporting mass from the distribution $\mathbb{P}(D_A)$ to $\mathbb{P}(D_B)$.

However, the kernel based on the squared Wasserstein distance, i.e., $k(A, B) \propto e^{-W_2^2(\mathbb{P}(D_A), \mathbb{P}(D_B))}$ may not be PSD when the data dimension exceeds 1 [2, 22], thus violating **(I)**. Moreover, as computing the Wasserstein distance involves an optimization problem, the most efficient method [7] still takes $\tilde{O}(\max(|D_A|, |D_B|)^2)$ time.³ This computation burden becomes expensive when repeated $O(a^2 + ab + b^2)$ times for each kernel entry, violating **(II)**. Thus, we must use alternatives to the Wasserstein distance, such as the *sliced Wasserstein distance* (SW) [22] that provably satisfies **(I)**.

PROPOSITION 1 ([22]). *The exponential kernels $k(A, B)$ based on the SW distance, including $\exp(-\gamma \cdot \text{SW}_2^{2\rho}(\mathbb{P}(D_A), \mathbb{P}(D_B)))$ and $\exp(-\gamma \cdot \text{SW}_1^\rho(\mathbb{P}(D_A), \mathbb{P}(D_B)))$ are positive semi-definite (PSD) and valid for $\gamma > 0$ and $\rho \in [0, 1]$.*

To address **(II)**, the SW distance can be efficiently approximated using Monte Carlo sampling with L projections. After each projection, the Wasserstein distance between one-dimensional distributions can be computed analytically. We additionally observe that only line 10 in Algorithm 1, which merges two sorted projection lists, is unique to the coalition pair (A, B) . Thus, only this step, which takes $O(L \cdot (|D_A| + |D_B|))$ time, is repeated for the $O(a^2 + ab + b^2)$ unique kernel entries. The L factor can be further reduced by parallelizing the computation for multiple projected directions. Given D_A has dimension m , steps 1-9 in Algorithm 1 only need to be precomputed once in $O(L|D_A|m + L|D_A| \log |D_A|)$ time.

³ \tilde{O} hides polylogarithmic factors.

Algorithm 1 Sliced Wasserstein Distance Computation

Input: Two dataset matrices X_A, X_B with m features/columns

Parameter: Number of projected directions L

Output: Sliced Wasserstein Distance $\text{SW}(X_A, X_B)$ between X_A and X_B .

```

1:  $s \leftarrow 0$ 
2: for  $l = 1$  to  $L$  do
3:   Uniformly sample  $\theta^{(l)}$  distributed on unit sphere
4:    $\Pi_{\theta^{(l)}}(X_A) \leftarrow X_A \theta^{(l)}$ 
5:    $\Pi_{\theta^{(l)}}(X_A) \leftarrow \text{sort}(\Pi_{\theta^{(l)}}(X_A))$ 
6:    $\Pi_{\theta^{(l)}}(X_B) \leftarrow X_B \theta^{(l)}$ 
7:    $\Pi_{\theta^{(l)}}(X_B) \leftarrow \text{sort}(\Pi_{\theta^{(l)}}(X_B))$ 
8:    $F_{X_A} \leftarrow$  the empirical c.d.f. with  $\Pi_{\theta^{(l)}}(X_A)$ 
9:    $F_{X_B} \leftarrow$  the empirical c.d.f. with  $\Pi_{\theta^{(l)}}(X_B)$ 
10:   $s \leftarrow s + \frac{1}{L} \left( \int_0^1 |F_{X_A}^{-1}(z) - F_{X_B}^{-1}(z)|^p dz \right)^{1/p}$ 
11: end for
12: return  $s$ 
```

2.2 Supervised Learning

In this section, we will additionally consider the target outputs y_i for each owner i . Specifically, we will define a transformation \mathcal{G}_η that will map the dataset (consists of both the input matrix and target outputs) to a common feature space.

DEFINITION 1. *Given the transformation \mathcal{G}_η , the supervised sliced Wasserstein (SSW) distance between the empirical data distribution D_A and D_B is*

$$\text{SSW}_2^{2\rho}(\mathbb{P}(D_A), \mathbb{P}(D_B); \mathcal{G}_\eta) = \text{SW}_2^{2\rho}(\mathbb{P}(\mathcal{G}_\eta(D_A)), \mathbb{P}(\mathcal{G}_\eta(D_B))) .$$

Supervised regression. Data valuation can be applied on regression problems, where the output y_i is a vector of real values. We concatenate y_i with X_i and vary the weight on y_i using a parameter η . Formally, let $\mathcal{G}_\eta(D_C) = \eta X_C \oplus (1 - \eta) y_C$, where $\eta \in (0, 1]$ is the scaling weight for the feature space.

Supervised classification. Data valuation is also often applied in supervised classification problems, where the output y_i is a vector of discrete class labels. At first glance, we can concatenate y_i with X_i , however, how do we measure the distance between different labels such as ‘airplane’, ‘bird’, and ‘truck’ in CIFAR-10? Inspired by Alvarez-Melis and Fusi [1], we quantify the distance between labels y^k and y^j as the sliced Wasserstein distance between the aggregated datasets with corresponding labels y^k and y^j , i.e., $\text{SW}(\mathbb{P}(x_i \mid (x_i, y^k) \in D_N), \mathbb{P}(x_i \mid (x_i, y^j) \in D_N))$. Thus, the ‘bird’ label is closer to ‘airplane’ than ‘truck’. We then use multi-dimensional scaling (MDS) [5] to embed each class y^j as a vector $e(y^j)$, preserving the distances between class labels. The function e is also applied element-wise to embed the target outputs y_C . Formally, let $\mathcal{G}_\eta(D_C) = \eta X_C \oplus (1 - \eta)e(y_C)$, where $\eta \in (0, 1]$ are scaling weights for the feature and label spaces, respectively.

PROPOSITION 2. *The kernel $k(A, B) = e^{-\gamma \text{SSW}_2^{2\rho}(\mathbb{P}(D_A), \mathbb{P}(D_B); \mathcal{G}_\eta)}$ is a valid kernel when $\gamma > 0$ and $\rho \in [0, 1]$ and $\eta > 0$.*

The proof based on Proposition 1 can be found in Appendix B.2. The following toy example demonstrates why our chosen kernel makes intuitive sense. In Figure 2, as owners l and l' own similar

data, the similarity $k(C \cup l, C \cup l')$ is large for any coalition C . In contrast, as owners l and j have very different data, the similarity $k(C \cup l, C \cup j)$ is always small. Thus, in a GP model, an observation of the utility of the coalition with l would greatly reduce the uncertainty of that with l' but not j .

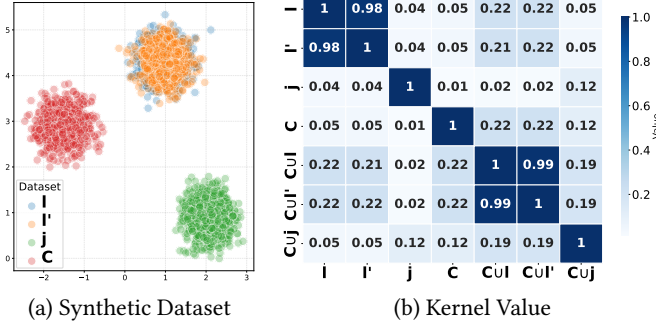


Figure 2: Datasets l and l' are similar while datasets l and j are very different. The SSW kernel similarity $k(C \cup l, C \cup l') = .99$ is large as compared to the similarity $k(C \cup l, C \cup j) = 0.19$.

3 SEMIVALUE ESTIMATION

We have defined the GP model and kernel to predict utilities for coalitions in \mathcal{B} . Given the a actual utility evaluations of \mathcal{A} (i.e., $\mathbf{u}_{\mathcal{A}}$), we first learn the kernel hyperparameters (e.g., γ) and choose the norm of the SW distance and scaling weights η for the feature and output spaces to maximize the log-likelihood. Then, we predict the posterior belief of the utilities $\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}$ of coalitions in \mathcal{B} . How do we compute the semivalue, such as the Shapley value, based on the observed utilities $\mathbf{u}_{\mathcal{A}}$ of coalitions in \mathcal{A} and the posterior belief $\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}$ of coalitions in \mathcal{B} ?

PROPOSITION 3 (SEMIVALUE PREDICTION). Let $\mathbf{w}_{\mathcal{A}}^i$ and $\mathbf{w}_{\mathcal{B}|\mathcal{A}}^i$ denote the vectors containing the weights of all coalitions in \mathcal{A} and \mathcal{B} respectively. The j -th entry of $\mathbf{w}_{\mathcal{A}}^i$ corresponds to the weight of coalition A_j in the computation of (i) the Shapley value $\phi_u(i)$ and (ii) the semivalue $\varphi_u(i)$, as defined in Equations 1 and 2). When $i \notin A_j$, the j -th entry is given by (i) $-1/(n - |A_j|)$ or (ii) $-\omega_{|A_j|}$. When $i \in A_j$, the j -th entry is (i) $1/(n - |A_j| - 1)$ or (ii) $\omega_{|A_j| - 1}$.

The estimated Shapley value and semivalue for owner i , denoted as $\hat{\phi}_i$ and $\hat{\varphi}_i$, are given by the weighted sum $\mathbf{w}_{\mathcal{A}}^i \top \mathbf{u}_{\mathcal{A}} + \mathbf{w}_{\mathcal{B}|\mathcal{A}}^i \top \hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}$ which follows the distribution

$$\mathcal{N}(\mathbf{w}_{\mathcal{A}}^i \top \mathbf{u}_{\mathcal{A}} + \mathbf{w}_{\mathcal{B}|\mathcal{A}}^i \top \mathbb{E}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}], \mathbf{w}_{\mathcal{B}|\mathcal{A}}^i \top \mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}] \mathbf{w}_{\mathcal{B}|\mathcal{A}}^i).$$

REMARK 1 (MONTE CARLO APPROXIMATION). The weights $\mathbf{w}_{\mathcal{A}}^i$ and $\mathbf{w}_{\mathcal{B}|\mathcal{A}}^i$ used to estimate $\hat{\phi}_i$ can also correspond to the weights used in Monte Carlo estimates of Shapley value [9, 16], ϕ_i' . In practice, when using Monte Carlo estimates, the sampled coalitions are collected and subsequently partitioned for actual evaluation and prediction, as illustrated in Figure 1.

Proposition 3 considers only the uncertainty arising from the use of GP model predictions instead of actual utility evaluations

(refer to (ii) in Section 1). When Monte Carlo approximation is used, the total uncertainty in the estimate $\hat{\phi}_i'$ should also consider the uncertainty introduced by Monte Carlo sampling. Formally, let σ_{MC} denote the standard deviation introduced by Monte Carlo methods, such as those described by Kolpaczki et al. [16], which evaluates and utilizes only a subset of all coalitions (see (i) in Section 1). Let σ_{GP} represent the standard deviation associated with our GP model, defined as the square root of the variance in Proposition 3. The total uncertainty $\mathbb{V}[\hat{\phi}_i']$ is upper bounded by $\sigma_{GP}^2 + \sigma_{MC}^2 + 2\sigma_{GP}\sigma_{MC}$.

3.1 Active Querying to Accelerate Uncertainty Reduction

We can reduce the variance σ_{GP}^2 associated with our GP model further by additionally evaluating \bar{b} coalitions $\bar{\mathcal{B}} \subset \mathcal{B}$. In particular, for each $B \in \bar{\mathcal{B}}$, the mediator trains a model on the aggregated data D_B to evaluate $u(B)$ and only predicts the utility of the remaining coalitions $\mathcal{B} \setminus \bar{\mathcal{B}}$. Instead of randomly selecting \bar{b} coalitions, the mediator can actively select the \bar{b} coalitions that lead to the largest reduction in the semivalue variance:

$$\bar{\mathcal{B}} = \operatorname{argmax}_{\mathcal{C} \subseteq \mathcal{B}, |\mathcal{C}|=\bar{b}} \mathbf{w}_{\mathcal{B}}^i \top \cdot (\mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}] - \mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A} \oplus \mathcal{C}}]) \cdot \mathbf{w}_{\mathcal{B}}^i. \quad (3)$$

Here, \oplus denote the concatenation operator and $\mathbb{V}[\hat{\mathbf{u}}_{\mathcal{B}|\mathcal{A}}]$ is the predictive variance of the utilities in \mathcal{B} given observed utilities from \mathcal{A} . As the weighted variance reduction function is often monotone submodular [4], Equation (3) can be maximized by the greedy algorithm in Algorithm 2.

Algorithm 2 Greedy Active Selection with Efficient Inverse Update

Input: Evaluated coalitions and utilities, $(\mathcal{A}, \mathbf{u}_{\mathcal{A}})$; unevaluated coalitions, \mathcal{B} ; kernel function k ; semivalue weight vector $\mathbf{w}_{\mathcal{B}}^i$

Parameter: Number of additional evaluations \bar{b}

Output: Selected coalitions $\bar{\mathcal{B}}$ for additional evaluations

- 1: Initialize: $\bar{\mathcal{B}} \leftarrow ()$, $\mathbf{K}^{-1} \leftarrow (\mathbf{K}_{\mathcal{A},\mathcal{A}} + \sigma^2 \mathbf{I})^{-1}$
- 2: **for** $j = 1$ **to** \bar{b} **do**
- 3: Set $\text{max_VR} \leftarrow 0$, $G^* \leftarrow \text{null}$
- 4: **for** $G \in \mathcal{B} \setminus \bar{\mathcal{B}}$ **do**
- 5: Let $C \leftarrow \bar{\mathcal{B}} \oplus (G,)$
- 6: Incrementally update \mathbf{K}^{-1} using the previous inverse and $k(G, \mathcal{A} \oplus \bar{\mathcal{B}})$ (as described in Appendix C.3)
- 7: Compute the variance reduction:

$$\text{VR} \leftarrow (\mathbf{w}_{\mathcal{B}}^i)^\top \cdot \mathbf{K}_{\mathcal{B},\mathcal{A} \oplus C} \mathbf{K}^{-1} \mathbf{K}_{\mathcal{A} \oplus C, \mathcal{B}} \cdot \mathbf{w}_{\mathcal{B}}^i$$
- 8: **if** $\text{VR} \geq \text{max_VR}$ **then**
- 9: Update $\text{max_VR} \leftarrow \text{VR}$, $G^* \leftarrow G$
- 10: **end if**
- 11: **end for**
- 12: Add the selected coalition: $\bar{\mathcal{B}} \leftarrow \bar{\mathcal{B}} \oplus (G^*,)$
- 13: **end for**
- 14: **return** $\bar{\mathcal{B}}$

4 EXPERIMENTS

We conduct experiments across several datasets to evaluate the effectiveness of our methods. For classification tasks, we utilize (a)

the Moon dataset [24], (b) the MNIST dataset [6], (c) the CIFAR-10 dataset [17], and (d) the IMDb dataset [20]. For regression tasks, we employ (e) the California Housing dataset (CaliH) [23], which provides real-world housing data. We use accuracy as the utility function for classification tasks and the R^2 score for regression tasks, defined as $R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}} = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$. Finally, in our experiments, we employ the Shapley value, the most widely used semivalue.

Our experiments aim to achieve three primary objectives: (i) to investigate how factors such as the kernel and the number of randomly and actively selected coalitions affect the performance and computation time of our framework in computing Shapley values; (ii) to demonstrate the application of our framework in data valuation, particularly for classification tasks; and (iii) to highlight the advantages of uncertainty quantification.

Baseline: We evaluate three baselines to compare with our proposed approach. The first baseline, denoted as OTDD, is based on the label-feature distance concept from Alvarez-Melis and Fusi [1]. Although this baseline employs the optimal transport distance between feature-label pairs, it is important to note that the exponential kernel derived from OTDD is not valid, unlike the kernel we propose. We summarize the computational complexity of each distance metric in Appendix B.3, Table 1. For this baseline, we implement the Sliced Wasserstein distance as the label distance for efficiency.⁴ The second and third baselines, inspired by Wang et al. [31], represent each dataset D_i with a binary indicator vector (01 encoding) $b_i \in \{0, 1\}^n$, indicating the indices present for each data owner i . These methods are referred to as GP-binary and NN-binary, employing Gaussian Process Regression and Neural Networks, respectively. Finally, for our SSW kernel, we set $\eta = 0.5$.

4.1 Coalition Utility Prediction

This section empirically evaluates the utility prediction performance of our proposed kernel SSW and the baseline methods. We compare the computation time between our DUPRE framework and an exhaustive evaluation of all possible coalitions. Our experiments include two settings: one where α randomly selected coalitions' utilities are evaluated by model training and another (with -a suffix) where half of the coalitions are randomly selected and the other $\alpha/2$ coalitions are actively selected as described in Algorithm 2 and Figure 1. After evaluating these α coalitions, we train our DUPRE framework on their utilities and then estimate the utilities of all remaining coalitions.

To assess the effectiveness of our methods, we compute the mean and standard deviation of the mean squared error (MSE) between our predictions and the actual utilities over ten runs, each using a different set of randomly evaluated coalitions (with random seeds from 0 to 9). We also calculate the Pearson correlation coefficient [14] to evaluate the correlation between predicted and actual utilities and use Kendall's tau metric [12] to assess the ranking order of predicted and actual Shapley values, which is critical for understanding a data owner's contribution.

⁴Alvarez-Melis and Fusi [1] originally defined the distance as $(\|x - x'\|_p + W_p^p)^{1/p}$, where W_p^p is the p -Wasserstein distance between label distributions. In our implementation, we use the Sliced Wasserstein (SW) distance for improved efficiency.

We first consider predicting the validation accuracy on the MNIST dataset, using a neural network as the ML model for evaluation. Each of the 6 data owners holds a distinct subset of digit labels: $\{1, 4, 5, 7, 8\}$, $\{2\}$, $\{9\}$, $\{6\}$, $\{0\}$, $\{3\}$. As shown in Figure 3, our kernel (SSW) outperforms the other baselines, evidenced by a lower mean squared error (MSE) in (a) and higher Pearson correlation coefficients in (c). While the active selection process improves performance, it also increases computation time. Nonetheless, our DUPRE framework remains faster than an exhaustive actual evaluation of all possible coalitions. Moreover, our kernel provides a superior ranking of Shapley values than other kernels, as indicated by the higher Kendall's tau coefficients in (d).

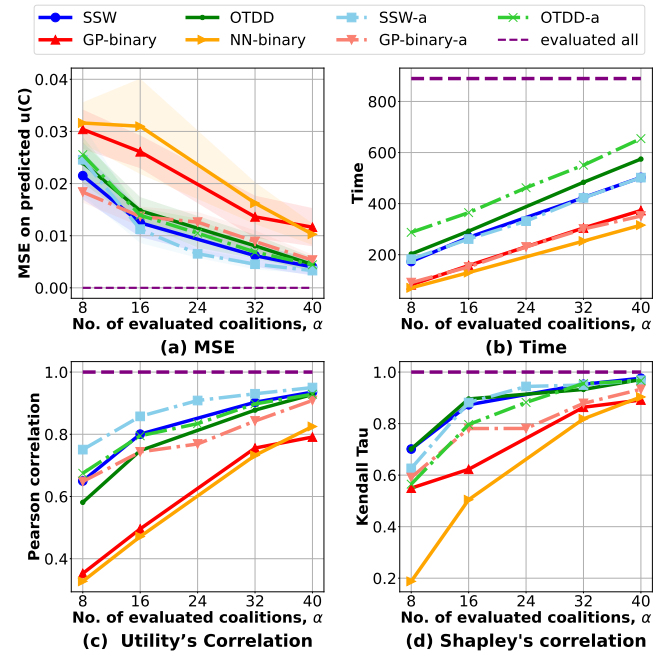


Figure 3: A comparison of the quality of the utility predictions and time taken for various methods on the MNIST dataset with 6 owners. The -a suffix indicates that we have actively selected 50% of the coalitions to accelerate uncertainty reduction.

We repeat this experiment on the regression dataset, CaliH, as illustrated in Figure 4, and use Multi-Layer Perceptron (MLP) as the ML model. The kernel based on OTDD is not applicable here, as it requires classification labels. Once again, our proposed kernel results in lower MSE and higher correlations than the baselines.

4.2 Evaluating the Quality of Shapley Value Predictions

In this section, we explore two approaches for computing the Shapley value: predicting the utilities of all coalitions to estimate exact Shapley values and predicting only the utilities of a subset to estimate approximate Shapley values.

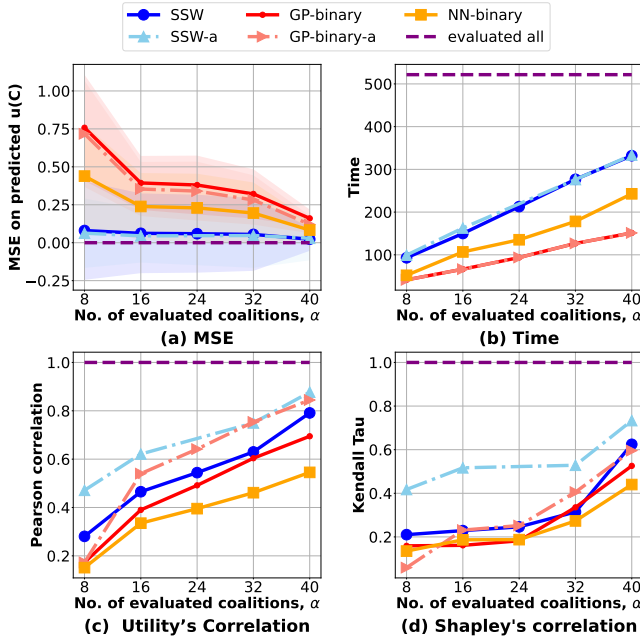


Figure 4: A comparison of the quality of the utility predictions and time taken for various methods on the CaliH dataset with 6 owners. The -a suffix indicates that we have actively selected 50% of the coalitions to accelerate uncertainty reduction.

4.2.1 Exact Shapley Value Estimation. We evaluated our framework on the CIFAR-10 and CaliH datasets with 8 data owners. For CIFAR-10, the utility function is the accuracy of the trained ResNet model. In Table 1, the first segment computes the Shapley value based on the predicted utilities from a GP model that is trained on the utility of 100 actually evaluated coalitions. In contrast, the GP model in the second segment is additionally trained on the utility of 10 more coalitions, randomly selected or selected by Algorithm 2.

We use Pearson and Kendall’s tau correlation coefficients as evaluation metrics to assess the agreement between our predicted and the exact Shapley values, i.e., $(\hat{\phi}_i)_{i \in N}$ and $(\phi_i)_{i \in N}$. We also compare against LAVA [11], a model-agnostic method that estimates Shapley values based on the distance between datasets and the task dataset, without requiring ML model training. Our method, SSW, produces estimates that are more correlated with the exact Shapley values, as evidenced by higher Pearson and Kendall’s tau coefficients as compared to other kernels. In contrast, the LAVA method results in the lowest correlations. Additionally, we observe that active selection, which incurs a slightly higher computation cost, improves the correlation more than random selection.

4.2.2 Approximate Shapley Value Estimation. Next, we evaluate our approach using the MNIST dataset distributed among 10 data owners. We compute the approximate Shapley values using permutation sampling, as outlined by Castro et al. [3]. We limit the number of permutation samples considered by the total number of

evaluated and predicted coalitions. Refer to Appendix C.4 for more details.

In our experiments, we actually evaluate the utilities of 512 coalitions by ML model training and use these utilities to train our GP model. Then, we either consider evaluating the utilities of another additional 250 coalitions (purple line) or predicting the utilities of β additional coalitions ($\beta \in [0, 250]$) using the GP model with different kernels. The key objective of the experiment is to determine if these predicted utilities can effectively substitute for actual utility evaluations.

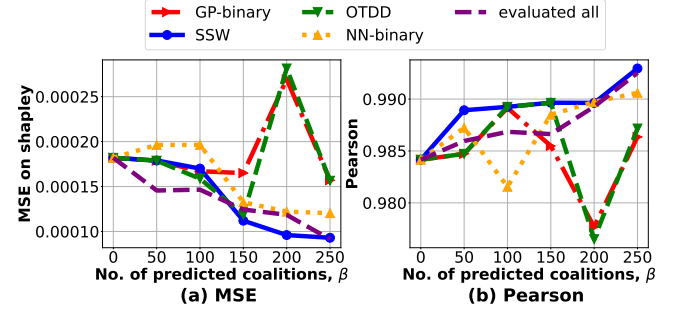


Figure 5: A comparison of the quality of the approximate Shapley value predictions for various methods on the MNIST dataset. The approximate Shapley value is computed using 512 actual utility evaluations and an additional β actual/predicted coalition utilities.

In Figure 5, the x-axis represents the number of additional coalitions used to compute the approximate Shapley value $\hat{\phi}_i$. We measure the mean squared error (MSE) between the estimated approximate Shapley value $\hat{\phi}_i$ and the exact Shapley value ϕ_i computed using Equation 1 as well as the Pearson correlation between the estimated approximate Shapley values and exact Shapley values across data owners. The purple lines exhibit a strictly decreasing MSE and a strictly increasing Pearson correlation as more coalitions are evaluated. Similarly, for our SSW kernel, both the MSE and Pearson correlation improve as more coalitions are predicted. Given that the performance of the SSW kernel mirrors that of the actual evaluations, it is a suitable substitute. In contrast, the other baselines do not exhibit the same trend; for instance, with NN-binary, increasing the number of predicted utilities can worsen both the MSE and Pearson correlation.

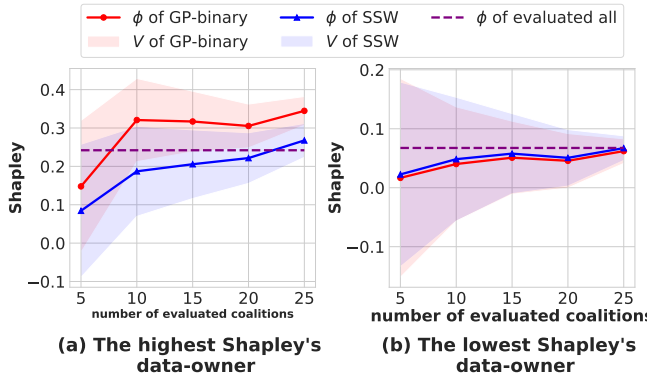
4.3 Benefits of Uncertainty Quantification

In this experiment, we examine the uncertainty of the predicted Shapley values of the MNIST classification task where the dataset is split among 5 data owners. We specifically focus on the data owner with (a) the highest contribution and (b) the lowest contribution. Since our kernel SSW outperforms OTDD, we consider only GP-binary as the baseline. We compute the Shapley values and their variances using the formula in Section 2. The result is illustrated in Figure 6.

As the number of evaluated coalitions increases, the predicted Shapley value gets closer to the actual Shapley value and the variance decreases. For our SSW kernel, the actual Shapley values always

Table 1: A comparison of the quality of exact Shapley value approximation and the time taken on the CIFAR-10 and CaliH datasets with 8 data owners. A higher correlation is preferred.

Method	CIFAR-10			CaliH		
	Pearson	Kendall Tau	Time (s)	Pearson	Kendall Tau	Time (s)
SSW	0.901 \pm 0.07	0.664 \pm 0.152	5893 \pm 1375	0.775 \pm 0.181	0.6714 \pm 0.18	543 \pm 10
OTDD	0.640 \pm 0.09	0.523 \pm 0.21	8056 \pm 2141	-	-	-
GP-binary	0.785 \pm 0.006	0.565 \pm 0.1	4120 \pm 593	0.528 \pm 0.134	0.593 \pm 0.129	264 \pm 35
NN-binary	0.612 \pm 0.01	0.544 \pm 0.12	2541 \pm 256	0.579 \pm 0.154	0.602 \pm 0.163	351 \pm 45
LAVA	-0.0785 \pm 0.0405	-0.3045 \pm 0.03	5580 \pm 394	0.1644 \pm 0.226	0.107 \pm 0.15	50 \pm 10
Evaluate 10 additional coalitions						
SSW - random	0.911 \pm 0.04	0.674 \pm 0.132	6137 \pm 1098	0.805 \pm 0.155	0.7124 \pm 0.21	585 \pm 20
SSW - active	0.934 \pm 0.07	0.677 \pm 0.126	6317 \pm 1567	0.831 \pm 0.165	0.7624 \pm 0.24	627 \pm 23
GP-binary - active	0.855 \pm 0.006	0.615 \pm 0.15	4320 \pm 635	0.655 \pm 0.1	0.653 \pm 0.120	388 \pm 40
Evaluate all coalitions	-	-	12458	-	-	950

**Figure 6: Plot of the GP model's predictive mean and standard deviation (shaded region around the line) of Shapley values for five data owners.**

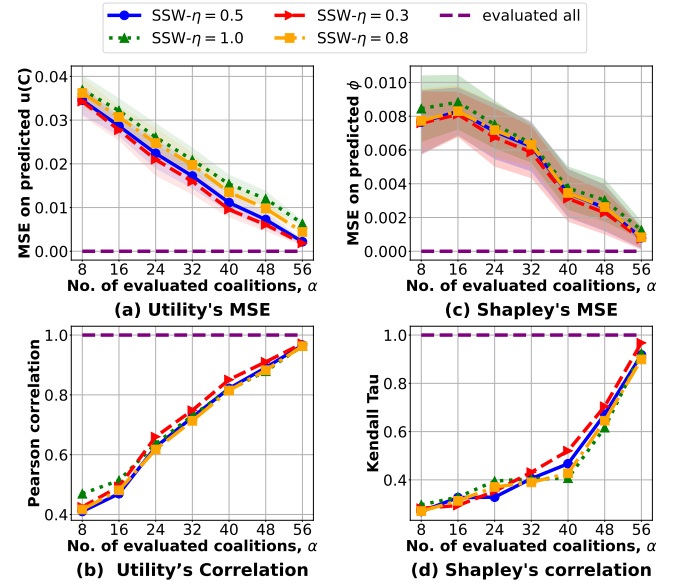
lie within the shaded region, suggesting our uncertainty quantification is well-calibrated.

4.4 Further Analysis

In this section, we present additional experiments to analyze and stress-test our framework. First, we perform an ablation study on the parameter η (Section 4.4.1), which controls the relative importance of label information in our kernel. Next, we demonstrate how to handle more complex datasets like IMDB (Section 4.4.2). Finally, we show the robustness of our method in a heavily imbalanced and heterogeneous setting using the IMDB dataset (Section 4.4.3).

4.4.1 Effect of the label-weight parameter η . We perform an ablation study to understand the influence of different η (i.e., different weight of the output label) affects the mean squared error (MSE) of the utility and Shapley value predictions. This experiment uses a classification task on the synthetic Moon dataset with 6 data owners.

A smaller η value means that the label information has a greater influence on the dataset distance, allowing the GP model to better capture label-dependent patterns. As illustrated in Figure 7, $\eta = 0.3$

**Figure 7: A comparison of the quality of the utility and Shapley value predictions for various η values on the Moon dataset with 6 data-owners.**

results in the lowest MSE and highest Pearson correlation, indicating that assigning more weight to label information leads to improved predictive performance.

4.4.2 Evaluation on an unstructured dataset with 10 data owners. We now evaluate our framework on the IMDB dataset, which comprises 50000 movie reviews labeled as either positive or negative. Following the OpenDataVal benchmark [10], we use DistilBERT [25] embeddings for each review. We split the dataset among 10 data owners and consider two experimental settings as in Sections 4.1 and 4.2.2). For the former, we use 256 evaluated coalitions. For the latter, we use 512 evaluated coalitions and 100 predicted coalitions.

Table 2 shows that our method SSW always achieves lower MSE and higher correlation than GP-binary and NN-binary, across five

Table 2: A comparison of the quality of utility predictions and Shapley value predictions for various methods on the IMDB dataset with 10 owners. Setup 1 (similar to Section 4.1) compares the actual utility $u(C)$ with GP predicted expected utility $\hat{u}(C)$ of various coalition C . Setup 2 (similar to Section 4.2.2) compares the actual Shapley value ϕ_i and the predicted Shapley value ($\hat{\phi}_i$) for various owner i . The results are the mean \pm std. over 5 runs. Lower MSE and higher correlation are preferred.

Method	Setup 1 (Section 4.1)			Setup 2 (Section 4.2.2)	
	$\text{MSE}(u(C), \hat{u}(C))$	$\text{Pearson}(u(C), \hat{u}(C))$	Shapley Corr.	$\text{MSE}((\phi_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$	$\text{Pearson}((\phi_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$
SSW (Ours)	$8.6 \times 10^{-6} \pm 0.008$	0.60 ± 0.13	0.75 ± 0.21	0.00022 ± 0.00016	0.59 ± 0.22
GP-binary	$1.8 \times 10^{-5} \pm 0.0042$	0.466 ± 0.17	0.66 ± 0.26	0.00029 ± 0.00014	0.50 ± 0.24
NN-binary	0.006 ± 0.008	-0.199 ± 0.07	-0.137 ± 0.35	0.0001 ± 0.00015	0.52 ± 0.025

runs where the evaluated and predicted coalitions are randomly varied. In Appendix D.2, we further validate our framework on the *Stanford Sentiment Treebank (SST-2)* [28] dataset. We also observe that SSW consistently outperforms the baselines across different text-based tasks and pre-trained embeddings.

4.4.3 Robustness to Heterogeneous Data Size and Distribution. While the earlier experiments already account for varying dataset sizes and distributions (see Appendix D, Table 2), we further validate our framework under more extreme heterogeneity using the IMDB dataset. Specifically, we consider 20 data owners, where the j -th data owner has $100j$ data points. Additionally, owners $j = 1 \dots 5$ hold only data of the positive class, while owners $j = 6 \dots 10$ hold only data of the negative class. The remaining 29000 data points are used as the validation set.

Table 3: A comparison of the quality of the Shapley value predictions for various methods on the IMDB dataset with 20 heterogeneous data owners. The Shapley value are computed based on 5000 actual utility evaluations and 20000 predicted utility evaluations.

Method	MSE	Pearson
SSW (Ours)	9.60×10^{-5}	0.736
GP-binary	9.48×10^{-5}	0.710
NN-binary	1.00×10^{-4}	0.652

We consider estimating the Shapley value with 25000 coalitions (~ 1800 permutations, see Appendix C.4). Out of these 25000 coalitions, we actually evaluate the utilities of 5000 random coalitions and predict the utilities of the remaining coalitions. In Table 3, we observe that despite the large variation in dataset sizes and distributions, SSW achieves the strongest correlation (0.736) and low MSE (9.60×10^{-5}). Additionally, we observe that the MSE does not increase when the owners have more data points. The Pearson correlation between dataset size and MSE for SSW is low and only 0.241.

5 RELATED WORK

Our work is complementary to related works on *data valuation* that propose new data utility functions (such as data volume [33] and information gain [27]) and strategies (such as the Shapley value [8, 18, 34], the Banzhaf value [30] and Least Core [34]). DUPRE can

be used to efficiently predict the utilities of any data utility function needed in the data valuation strategies.

Our work is also complementary to *semivalue approximation techniques* that reduce the number of coalitions to evaluate such as permutation sampling [3], stratified sampling [21], structured sampling [29], or approximating Shapley without marginal contribution [16]. Instead of evaluating all the sampled coalitions' utilities by training a model, we propose evaluating a subset and predicting the remaining utilities using a GP model. Our work offers an alternative to methods that reduce the cost per evaluation, such as TMC-Shapley, Gradient Shapley [8] and the influence function heuristic used by Jia et al. [9].

Our work can be extended to make use of other *dataset distances*, such as optimal transport dataset distance (OTDD) [1], if they satisfy the valid properties of a kernel. There are also other data valuation works that have used the SW distance or aim to reduce the cost of each utility function but they differ in their application. In data valuation works, Just et al. [11], Kessler et al. [13] have also defined their data utility function based on the optimal transport distance between training data subsets and validation data. However, our purpose of considering dataset distances differs.

6 CONCLUSION

In this paper, we introduce DUPRE, a novel framework that complements existing sampling-based approximation methods to further boost the efficiency of computing CGT-based data valuation. We design a valid kernel based on the sliced Wasserstein distance and adapt the distance to consider the target outputs in supervised learning. As our kernel can encode prior knowledge of similarities between different data subsets, our GP model outperforms other approaches in our experiments.

While DUPRE demonstrates strong empirical performance, its predictions are not guaranteed to be accurate for every dataset or utility function. We recommend using a validation set of coalitions and utilities to continually assess and improve its predictions. Future work can consider other applications of data utility prediction, such as in online data valuation scenarios where new data owners frequently join or leave the collaboration.

ACKNOWLEDGMENTS

This research is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Program (AISG Award No: AISG2-RP-2020-018).

REFERENCES

- [1] David Alvarez-Melis and Nicolo Fusi. 2020. Geometric dataset distances via optimal transport. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., Virtual, 21428–21439.
- [2] François Bachoc, Alexandra Suvorikova, David Ginsbourger, Jean-Michel Loubes, and Vladimir Spokoiny. 2020. Gaussian processes with multidimensional distribution inputs via optimal transport and Hilbertian embedding. *Electronic journal of statistics* 14, 2 (2020), 2742–2772.
- [3] Javier Castro, Daniel Gómez, and Juan Tejada. 2009. Polynomial calculation of the Shapley value based on sampling. *Computers & operations research* 36, 5 (2009), 1726–1730.
- [4] Abhimanyu Das and David Kempe. 2008. Algorithms for subset selection in linear regression. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (Victoria, British Columbia, Canada). Association for Computing Machinery, New York, NY, USA, 45–54.
- [5] Erik Demaine, Adam Hesterberg, Frederic Koehler, Jayson Lynch, and John Urschel. 2021. Multidimensional Scaling: Approximation and Complexity. In *Proceedings of the 38th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, Virtual, 2568–2578.
- [6] Li Deng. 2012. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine* 29, 6 (2012), 141–142.
- [7] Pavel Dvurechensky, Alexander Gasnikov, and Alexey Kroshnin. 2018. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. PMLR, Stockholm, Sweden, 1367–1376.
- [8] Amirata Ghorbani and James Zou. 2019. Data Shapley: Equitable Valuation of Data for Machine Learning. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. PMLR, CA, USA, 2242–2251.
- [9] 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 *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, Okinawa, Japan, 1167–1176.
- [10] Kevin Fu Jiang, Weixin Liang, James Zou, and Yongchan Kwon. 2023. Open-DataVal: a Unified Benchmark for Data Valuation. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, Vol. 36. Curran Associates Inc., Red Hook, NY, USA, 28624–28647.
- [11] Hoang Anh Just, Feiyang Kang, Tianhao Wang, Yi Zeng, Myeongseob Ko, Ming Jin, and Ruoxi Jia. 2023. LAVA: Data Valuation without Pre-Specified Learning Algorithms. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*. OpenReview.net, Kigali, Rwanda.
- [12] Maurice G. Kendall. 1938. A new measure of rank correlation. *Biometrika* 30, 1-2 (1938), 81–93.
- [13] Samuel Kessler, Tam Le, and Vu Nguyen. 2025. SAVA: Scalable Learning-Agnostic Data Valuation. In *The 13th International Conference on Learning Representations (ICLR)*. OpenReview.net, Singapore.
- [14] Wilhelm Kirch (Ed.). 2008. *Pearson’s Correlation Coefficient*. Springer Netherlands, Dordrecht, 1090–1091.
- [15] Pang Wei W Koh, Kai-Siang Ang, Hubert Teo, and Percy S Liang. 2019. On the Accuracy of Influence Functions for Measuring Group Effects. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc., Vancouver, BC, Canada.
- [16] Patrick Kolpaczki, Viktor Bengs, Maximilian Muschalik, and Eyke Hüllermeier. 2024. Approximating the Shapley value without marginal contributions. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, Vancouver, 13246–13255.
- [17] A. Krizhevsky and G. Hinton. 2009. *Learning multiple layers of features from tiny images*. Master’s thesis. Department of Computer Science, University of Toronto.
- [18] Yongchan Kwon and James Zou. 2022. Beta Shapley: a Unified and Noise-reduced Data Valuation Framework for Machine Learning. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, Valencia, Spain, 8780–8802.
- [19] Weida Li and Yaoliang Yu. 2024. Faster Approximation of Probabilistic and Distributional Values via Least Squares. In *Proceedings of the 20th International Conference on Learning Representations (ICLR)*. OpenReview.net, Vienna, Austria.
- [20] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Portland, Oregon, USA, 142–150.
- [21] Sasan Maleki, Long Tran-Thanh, Greg Hines, Talal Rahwan, and Alex Rogers. 2013. Bounding the estimation error of sampling-based shapley value approximation with/without stratifying. *CoRR, abs/1306.4265* 2, 1 (2013).
- [22] Dimitri Meunier, Massimiliano Pontil, and Carlo Ciliberto. 2022. Distribution regression with sliced Wasserstein kernels. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*. PMLR, Maryland, USA, 15501–15523.
- [23] Kelley Pace and Ronald Barry. 1997. Sparse spatial autoregressions. *Statistics & Probability Letters* 33, 3 (1997), 291–297.
- [24] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *The Journal of Machine Learning research* 12 (2011), 2825–2830.
- [25] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv:1910.01108
- [26] Rachael Hwee Ling Sim, Xinyi Xu, and Bryan Kian Hsiang Low. 2022. Data Valuation in Machine Learning: “Ingredients”, Strategies, and Open Challenges. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. International Joint Conferences on Artificial Intelligence Organization, Vienna, Austria, 5607–5614.
- [27] Rachael Hwee Ling Sim, Yehong Zhang, Mun Choon Chan, and Bryan Kian Hsiang Low. 2020. Collaborative machine learning with incentive-aware model rewards. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, Vienna, Austria, 8927–8936.
- [28] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Seattle, Washington, USA, 1631–1642.
- [29] Tjeerd van Campen, Herbert Hamers, Bart Husslage, and Roy Lindelauf. 2018. A new approximation method for the Shapley value applied to the WTC 9/11 terrorist attack. *Social Network Analysis and Mining* 8 (2018), 1–12.
- [30] Jiachen T Wang and Ruoxi Jia. 2023. Data Banzhaf: A robust data valuation framework for machine learning. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, Valencia, Spain, 6388–6421.
- [31] Tianhao Wang, Yu Yang, and Ruoxi Jia. 2022. Improving Cooperative Game Theory-based Data Valuation via Data Utility Learning. In *ICLR 2022 Workshop on Socially Responsible Machine Learning*. International Conference on Learning Representations, Virtual.
- [32] Christopher Williams and Carl Rasmussen. 1995. Gaussian processes for regression. In *Advances in neural information processing systems*. Vol. 8. MIT Press, CO, USA.
- [33] Xinyi Xu, Zhaoxuan Wu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. 2021. Validation free and replication robust volume-based data valuation. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc, Virtual, 10837–10848.
- [34] 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*. AAAI Press, Palo Alto, California, USA, 5751–5759.