

Data Valuation and Detections in Federated Learning

Wenqian Li
National University of Singapore
wenqian@u.nus.edu

Shuran Fu
NUS (Chongqing) Research Institution
shuran.fu@nusricq.cn

Fengrui, Zhang
Georgetown University
fz128@georgetown.edu

Yan Pang*
National University of Singapore
jamespang@nus.edu.sg

Abstract

Federated Learning (FL) enables collaborative model training while preserving the privacy of raw data. A challenge in this framework is the fair and efficient valuation of data, which is crucial for incentivizing clients to contribute high-quality data in the FL task. In scenarios involving numerous data clients within FL, it is often the case that only a subset of clients and datasets are pertinent to a specific learning task, while others might have either a negative or negligible impact on the model training process. This paper introduces a novel privacy-preserving method for evaluating client contributions and selecting relevant datasets without a pre-specified training algorithm in an FL task. Our proposed approach, FedBary, utilizes Wasserstein distance within the federated context, offering a new solution for data valuation in the FL framework. This method ensures transparent data valuation and efficient computation of the Wasserstein barycenter and reduces the dependence on validation datasets. Through extensive empirical experiments and theoretical analyses, we demonstrate the advantages of this data valuation method as a promising avenue for FL research. Codes are available at <https://github.com/muz1lee/MOTdata>.

1. Introduction

Federated Learning (FL) has emerged as a privacy-preserving approach for collaboratively training models without sharing data [21], in which the learning proceeds by iteratively exchanging model parameters between the server and clients. The success of the trained model hinges on the availability of large, high-quality, and rele-

vant data [13]. Thus, it is crucial for the server to select valuable clients in the FL training process to ensure better model performance and explainability [33]. In the context of cross-silo FL, there is a mutual interest among clients in understanding the value of their own data as well as that of others. If there are free riders or malicious actors in FL, clients are unwilling to provide high-quality data to train the local model. Thus, fair valuation of data quality and accurate detection of noisy data become indispensable components of federated training [20].

In the context of a data marketplace with numerous potential clients, each possessing unique datasets, a central server issuing a learning task must discern and select the most valuable data clients for participation. This necessitates a fair and accurate assessment of each client's contribution. Commonly, existing methods evaluate contributions based on validation performance of the trained model, utilizing the concept of Shapley value (SV) to measure each client's marginal contribution to an FL model. This method, however, involves evaluating all possible subset combinations, leading to a computational complexity of $\mathcal{O}(2^N)$ for N clients. Although efforts have been made to reduce the complexity through approximations [8, 11, 20, 34], these approaches may introduce biased evaluations and remain impractical for large-scale settings. Moreover, approaches based on validation performance are post-hoc, which assess client contributions after model training. This approach is problematic because some clients might be irrelevant to the FL task, and subsets of noisy data might also be involved in training without detection, leading to the waste of computational resources and poor model performance. Additionally, this evaluation method is typically client-level, based on model gradients, and does not delve into the granularity of individual data points, leading to a lack of transparency in the evaluation process.

*Corresponding Author

Consequently, alternative strategies have been explored for evaluating clients and their data prior to model training. For instance, [35] proposed a pre-training approach where small sizes of data samples are shared to achieve this. Recently, Wasserstein distance [2, 13, 14] have been introduced to evaluate data without specifying a learning algorithm in advance. However, these approaches require access to data samples, which may not be feasible in privacy-sensitive settings. Moreover, they also rely on the validation dataset for assessment. Given challenges inherent in previous works, our research is driven by following questions:

1) How to evaluate and select valuable data points without sharing any raw data? Previous methods can reasonably assess clients’ contributions, there is a lack of a well-established approach to achieving a comprehensive understanding of individual data point contributions in FL. This limitation hinders transparency and the persuasive power of evaluations. Furthermore, developing such an approach would help the server select the most valuable data for model training.

2) How to predict and evaluate data client contributions without involving model training? Previous approaches involve training federated models to evaluate validation performance or utilize generative models to learn data distributions [32]. These methods are often computationally intensive and reliant on validation data. Addressing this, the key challenge is to develop data valuation methods that do not require direct training of federated models. Such an advancement is not merely beneficial but crucial, particularly in FL scenarios characterized by complex model training processes or where validation datasets are not readily available on the FL server.

3) How to evaluate data contributors in a large-scale setting? As the number of clients increases, the complexity of previous methods based on Shapley value (SV) grows exponentially. Therefore, offering a lower-complexity data valuation method that does not require full data access will be fundamental for large-scale FL settings, especially those involving more than 100 clients.

Leveraging the advances in computational optimal transport [25] and recent developments in the federated scenario as in [28], we introduce FedBary as an innovative solution to address the aforementioned challenges. To the best of our knowledge, FedBary is the first established general privacy-preserving framework to evaluate both client contribution and individual data value within FL. We also provide an efficient algorithm for computing the Wasserstein barycenter, to lift dependence on validation data. FedBary offers a more transparent viewpoint for data evaluations in FL, which directly computes the distance among various data distributions, mitigating the fluctuations introduced by model training and testing. Furthermore, it demonstrates faster computations without sacrificing performance compared to tradi-

Methods	Privacy	Without validation	Without model	Noisy detection
DataSV [8]	✓	×	×	×
CGSV [42]	✓	✓	×	×
GTG-SV [20]	✓	×	×	×
Lava [13]	×	×	✓	✓
Ours (FedBary)	✓	✓	✓	✓

Table 1. Overview of different approaches: we aim to handle both client evaluation and data detection tasks, w/o validation data. We conduct extensive experiments and theoretical analysis to show the promising applications of this research. A comparison of representative state-of-the-art approaches is shown in Table 1.

2. Related Work

Optimal Transport Application Optimal Transport (OT) is a classical mathematical framework used to solve transportation and distribution problems [37, 38], and it has been used in the field of machine learning for different tasks, including model aggregation and domain adaption [7]. Its effectiveness has been proved empirically [4] and theoretically [9, 29]. [2] showed that it can be used to measure the distance between two datasets, providing a meaningful comparison of datasets and correlating well with transfer learning hardness. In order to get closer to a real situation, some research put their emphasis on multi-source domain adaption (MSDA), where there are multiple source domains and a robust model is required to perform well on any target mixture distribution [9]. The results of recent research showed that OT theory is capable of being used as a tool to solve multi-source domain adaption problems [23, 27, 30, 36, 40].

Data Valuation The topic of data quality valuation has become popular and gained research interest in recent years since the quality of data will have a direct impact on the trained models, thus influencing downstream tasks. The most popular underlying metric for data valuation algorithms is the Shapley value [8, 11, 18], which calculates the marginal contribution and measures the average change in the predefined function when a particular data point is removed. This function is commonly set as the performance of a model trained on a subset of the training dataset. Other papers following this method include KNNShapley [10], Volume-based Shapely [43], BetaShapley [16], DataBanzhaf [39] and AME [19]. Alternative underlying methods are also explored to evaluate data contribution. For example, gradient-based approaches like Influence Function [6] are proposed to quantify the rate of changes for a utility value when some data points are more weighted. [26] proposed *TracIn* method to estimate training data influence by tracing gradient descent. However, this method relies heavily on the training algorithms of machine learning models. [13] developed a proxy for the validation performance, and

their method Lava can be used to evaluate data in a way that is oblivious to the learning algorithms. Reinforcement learning [44] and out-of-bag estimation [17] are also introduced as tools to evaluate data values. A benchmark for data valuation is provided in [12], which can be conveniently accessed by data users and publishers.

3. Technical Preliminary

3.1. Wasserstein Distance and Barycenter

We denote by $\mathcal{P}(X)$ the set of probability measures in X and $\mathcal{P}_p(X)$ the subset of measures in $\mathcal{P}(X)$ with finite p -moment $p \geq 1$. For $P \in \mathcal{P}_p(X)$ and $Q \in \mathcal{P}_p(Y)$, with distance function $d(x, y)$, the p -Wasserstein distance $\mathcal{W}_p(P, Q)$ between the measures P and Q is defined as

$$\mathcal{W}_p(P, Q) = \left(\inf_{\pi \in \Pi(P, Q)} \int_{X \times Y} d^p(x, y) d\pi(x, y) \right)^{1/p} \quad (1)$$

here $\Pi(P, Q)$ denotes the set of all joint distributions on (X, Y) that are marginally distributed as P and Q , and $\mathcal{W}_p(P, Q)$ quantifies the optimal expected cost of mapping samples from P to Q . When the infimum in (1) is attained, any probability π that realizes the minimum is an optimal transport plan. The notion of Wasserstein barycenter can be viewed as the mean of probability distributions $P_i, i \in [1, N]$ in the Wasserstein space [1], defined as

Definition 1 Given P_i with $P_i \in \mathcal{P}_p(X), \forall i \in [1, N]$, with positive constants $\{\lambda_i\}_{i=1}^N$ such that $\sum_{i=1}^N \lambda_i = 1$, the Wasserstein barycenter is the optimal Q^* satisfying

$$Q^* = \arg \min \sum_{i=1}^N \lambda_i \mathcal{W}_p(P_i, Q). \quad (2)$$

The problem of calculating the Wasserstein barycenter of empirical measures \hat{P}_i is proposed by [5] since we can only access samples following distributions of corresponding P_i in common case. In addition, consider the infimal convolution cost in the multi-source optimal transport theory [24], p -Wasserstein distance with an N -ary distance function $d(x_1, \dots, x_N)$ among all distributions P_i is equivalent to find the minimal of the total pairwise between Q and P_i , that is $\mathcal{W}_p(P_1, \dots, P_n) = \min_Q \sum_{i=1}^N \mathcal{W}_p(Q, P_i)$, which provides a viewpoint of the mixture distribution of different source distributions and help measure the heterogeneity among them.

3.2. Geodesics and Interpolating Measure

This part hinges on the geometry of the Wasserstein distance and geodesics for applying in a federated manner [28]. We utilized the triangle inequality property of Wasserstein distance and defined the concept of geodesics and the interpolating measure. We show the details of the definitions in Appendix 7.

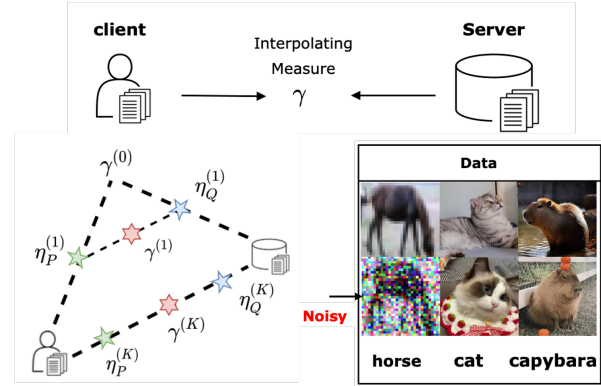


Figure 1. Client holds P and server holds Q , the interpolating measure γ aids to measure distance $\mathcal{W}_p(P, Q)$. Local interpolating measures η_p and η_Q are shared for calculation and detection.

Denote $\bar{\mu}$ as the geodesics, then any point γ_t on $\bar{\mu}$ is an interpolating measure between distribution P and Q , thus formulating the equality:

$$\mathcal{W}_p(P, Q) = \mathcal{W}_p(P, \gamma_t) + \mathcal{W}_p(\gamma_t, Q). \quad (3)$$

This provides insight for computing the Wasserstein distance in a federated manner: once we can approximate the interpolating measure between two distributions, we can measure the Wasserstein distance based on (3).

4. Problem Formulation

Suppose there are N benign clients and each client i holds dataset $\mathcal{D}_i = \{(x_{i,j}, y_{i,j})\}_{j=1}^m$ with size m . Considering data heterogeneity, we assume \mathcal{D}_i is independently and identically sampled from the different distribution P_i , but shares the same feature space \mathcal{X} and label space \mathcal{Y} such that $\mathcal{X} \times \mathcal{Y} = \mathcal{Z}$. Since we could not access the true distribution, in general, one can construct discrete measures $P_i(x, y) := \frac{1}{m} \sum_{j=1}^m \delta_{(x_{i,j}, y_{i,j})}$, where δ is a Dirac function. Clients are supposed to collaboratively train a model, and before training, the server wants to measure the contribution and select relevant data to a target distribution Q . If the server holds validation dataset $\mathcal{D}_Q = \{(x_{q,j}, y_{q,j})\}_{j=1}^m$, we assume is i.i.d to Q . If the server does not have a validation dataset, the target distribution is approximated by \hat{Q} . We use the Euclidean distance $d(\cdot, \cdot)$ to measure feature distance. For the label distance, we use the conditional feature space as $P_i(x|y_o) = \frac{P_i(x)\mathbb{I}[y=y_o]}{\int P_i(x)\mathbb{I}[y=y_o]dx}$ with $y_o \in \mathcal{Y}$ [2]. In order to calculate the Wasserstein distance $\mathcal{W}_p(P_i, Q)$, for each client, there is an interpolating measure γ_i to be approximated. Therefore, γ_i will be initialized and shared between the server and client i for iteratively updating. Any raw dataset D_Q and $D_i, \forall i \in [1, N]$ will not be shared.

In the following sections, we will first provide the technique to approximate Wasserstein barycenter in Section 4.1 for a better understanding of the evaluating procedure. Then, we will investigate the application of Wasserstein measure in scenarios with and without validation datasets

in Section 4.2. Furthermore, we leverage the duality theorem to detect noisy and irrelevant data points in Section 4.3. Figure 1 shows the overall framework.

4.1. Federated Wasserstein Barycenter

Our goal is to approximate the Wasserstein barycenter in Definition 1 among data distributions P_1, P_2, \dots, P_N on the server. Without loss of generality, we assume $\lambda_i = 1$. In order to make the update γ_i , the triangle inequality defined in Property 1 is extended in the following way,

$$\begin{aligned} \mathcal{W}_p(P_i, Q) &\leq \mathcal{W}_p(P_i, \gamma_i) + \mathcal{W}_p(\gamma_i, Q), \\ \mathcal{W}_p(P_i, \gamma_i) &= \mathcal{W}_p(P_i, \eta_{P_i}) + \mathcal{W}_p(\eta_{P_i}, \gamma_i), \\ \mathcal{W}_p(\gamma_i, Q) &= \mathcal{W}_p(\gamma_i, \eta_{Q_i}) + \mathcal{W}_p(\eta_{Q_i}, Q), \end{aligned} \quad (4)$$

where η_{P_i} is the interpolating measure between P_i and γ_i computed by i -th client, and η_{Q_i} is the interpolating measure between γ_i and Q computed by the server. The interpolating measure between P_i and γ_i (same way to Q and γ_i) could be approximated [28] based on

$$\eta_{P_i} = \frac{1}{m} \sum_{j=1}^m \delta_{(1-t)x_j + tm(\pi^* \mathbf{Z}')_j}, \quad (5)$$

where π^* is the optimal transportation plan between P_i and γ_i , x_j is j -th sample from P_i , and \mathbf{Z}' is the matrix of samples from γ_i . $t \in [0, 1]$ is a hyperparameter.

Once γ_i is an interpolating measure between P_i and Q , $\mathcal{W}_p(P_i, Q)$ equals the summation of four terms on the right-hand side in (4) and could be further applied to approximate the barycenter when $\min \sum_{i=1}^N \mathcal{W}_p(P_i, Q)$ is attained. Therefore, we develop a K -round iterative optimization procedure to approximate the interpolating measure $\gamma_i^{(k)}$ and the Wasserstein barycenter $Q^{(k)}$ based on $\eta_{P_i}^{(k)}$ and $\eta_{Q_i}^{(k)}$. Specifically, at each iteration k , the clients receive current iterate $\gamma_i^{(k-1)}$ and compute interpolating measure $\eta_{P_i}^{(k)}$ between P_i and $\gamma_i^{(k-1)}$, and server computes interpolating measure $\eta_{Q_i}^{(k)}$ between each $\gamma_i^{(k-1)}$ and $Q^{(k-1)}$. Then the server sends all $\eta_{Q_i}^{(k)}$ to clients, and i -th client computes the next iterate $\gamma_i^{(k)}$ based on $\eta_{P_i}^{(k)}$ and $\eta_{Q_i}^{(k)}$. The distance in an iterative procedure is calculated based on (6) as follows,

$$\begin{aligned} \mathcal{W}_p(P_i, Q) &\leq \mathcal{W}_p(P_i, \eta_{P_i}^{(k)}) + \mathcal{W}_p(\eta_{P_i}^{(k)}, \gamma_i^{(k-1)}) \\ &\quad + \mathcal{W}_p(\gamma_i^{(k-1)}, \eta_{Q_i}^{(k)}) + \mathcal{W}_p(\eta_{Q_i}^{(k)}, Q^{(k-1)}), \end{aligned} \quad (6)$$

where $\gamma_i^{(k)}$ is updated as [28] by

$$\gamma_i^{(k)} \in \operatorname{argmin} [\mathcal{W}_p(\eta_{P_i}^{(k)}, \gamma_i^{(k-1)}) + \mathcal{W}_p(\gamma_i^{(k-1)}, \eta_{Q_i}^{(k)})]. \quad (7)$$

It is straightforward to see the distribution Q is only involved in $\mathcal{W}_p(Q^{(k-1)}, \eta_{Q_i}^{(k)})$ and thus we can easily update barycenter $Q^{(k)}$ simultaneously by transporting samples from $\gamma_i^{(k)}$ to the common distribution.

As for the pairwise $z = (x, y)$, it is not easy to approximate the interpolating measure with classification labels. Based on the insight of [2, 28], we utilize the point-wise notion of distance in $\mathcal{X} \times \mathcal{Y}$ as

$$\begin{aligned} d((x, y), (x', y')) &\triangleq (d(x, x') + \mathcal{W}_2^2(\alpha_y, \alpha_{y'}))^{1/2}, \\ \mathcal{W}_2^2(\alpha_y, \alpha_{y'}) &= \|m_y - m_{y'}\|_2^2 + \|\Sigma_y - \Sigma_{y'}\|_2^2, \end{aligned} \quad (8)$$

where $\alpha_{y'}$ is conditional feature distribution $P(x|y = y')$ that follows the Gaussian distribution with mean m_y and covariance Σ_y , we can construct the augmented representation of each dataset, that each pair (x, y) is a stacked vector $\tilde{x} := [x; m_y; \operatorname{vec}(\Sigma_y^{1/2})]$. Then with the stacked matrix $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{X}}'$, the data distance is calculated based on $\mathcal{W}_p(\tilde{\mathbf{X}}, \tilde{\mathbf{X}}') \leq \mathcal{W}_p(\tilde{\mathbf{X}}, \gamma) + \mathcal{W}_p(\gamma, \tilde{\mathbf{X}}')$. We use $\tilde{\mathbf{X}}_i$ and $\tilde{\mathbf{X}}_Q$ as the stacked vector for the samples from P_i and Q . Our algorithm is summarized in Algorithm 1.

4.2. Evaluate Client Contribution

When the server has a validation set $D_Q = \{x_{q,j}, y_{q,j}\}_{j=1}^m$, then it can easily measure $\mathcal{W}_p(\tilde{\mathbf{X}}_i, \tilde{\mathbf{X}}_Q)$ without the initialization of $Q^{(0)}$ and update on $\tilde{\mathbf{X}}_Q^{(k)}$ (line 7 in Algorithm 1). From [13], we know the validation loss of the trained model is bounded by the distance between the training data and the validation data. Thus, we can measure the contribution of i -th using the reverse of Wasserstein distance $\mathcal{W}_p(\tilde{\mathbf{X}}_i, \tilde{\mathbf{X}}_Q)$ without training a federated model. A smaller distance leads to better performance on the validation data D_Q and can be considered more valuable.

Algorithm 1 FedBary

Input: Local data distribution P_i with $\tilde{\mathbf{X}}_i$, initialisation of $Q^{(0)}$ with support $\tilde{\mathbf{X}}_Q^{(0)}$ and $\gamma_i^{(0)}$, $i = 1, \dots, N$. (No initialisation of $\tilde{\mathbf{X}}_Q^{(0)}$ with validation set, use fixed $\tilde{\mathbf{X}}_Q$)

- 1: **for** $k = 1$ to K **do**
- 2: Clients compute distance $\mathcal{W}_p(\tilde{\mathbf{X}}_i, \gamma_i^{(k-1)}) = \mathcal{W}_p(\tilde{\mathbf{X}}_i, \eta_{P_i}^{(k)}) + \mathcal{W}_p(\eta_{P_i}^{(k)}, \gamma_i^{(k-1)})$
- 3: Clients send $\mathcal{W}_p(\tilde{\mathbf{X}}_i, \gamma_i^{(k-1)})$ and $\gamma_i^{(k-1)}$ to server.
- 4: Server computes distance $\mathcal{W}_p(\tilde{\mathbf{X}}_Q^{(k-1)}, \gamma_i^{(k-1)}) = \mathcal{W}_p(\tilde{\mathbf{X}}_Q^{(k-1)}, \eta_{Q_i}^{(k)}) + \mathcal{W}_p(\eta_{Q_i}^{(k)}, \gamma_i^{(k-1)})$
- 5: Server sends $\eta_{Q_i}^{(k)}$ to i -th client
- 6: Client i updates $\gamma_i^{(k)}$ based on $\eta_{P_i}^{(k)}$ and $\eta_{Q_i}^{(k)}$.
- 7: Server updates $\tilde{\mathbf{X}}_Q^{(k)}$ (If without validation set)
- 8: **end for**

Output: $\tilde{\mathbf{X}}_Q^{(K)}, \gamma_i^{(K)}, \eta_{P_i}^{(K)}, \eta_{Q_i}^{(K)}, \mathcal{W}_p(\tilde{\mathbf{X}}_Q^{(K)}, \tilde{\mathbf{X}}_i)$

However, if there is no validation dataset D_Q , the initialization of $Q^{(0)}$ with $\tilde{\mathbf{X}}_Q^{(0)}$ should follow the same dimension of constructed matrix $\tilde{\mathbf{X}}_i$ using D_i . Some work [22, 31]

discuss that the target distribution (the model truly learns) is the mixture of local distributions, i.e., $\mathcal{U}_\lambda = \sum_{i=1}^N \lambda_i P_i$ for some $\lambda \in \Delta_N$, $\Delta_N := \{\mathbf{p} \in [0, +\infty)^N : \langle \mathbf{p}, \mathbf{1}_N \rangle = 1\}$, which is equivalent to the λ -weighted Euclidean barycenter of the distribution P_1, \dots, P_N . In this paper, instead of the Euclidean barycenter, we opt for the Wasserstein barycenter $\lambda \mathcal{W}_p(\mathcal{U}, P_i)$ as the target distribution to evaluate the distance for two main reasons. Firstly, the Wasserstein barycenter is superior in capturing complex geometric structures. Secondly, employing a similar approach to approximate the Federated Euclidean barycenter poses a risk of data exposure. This risk arises from the inherent simplicity of the weighted averaging process. Due to some fluctuations and the different sampling schemes in the various FL algorithms, coefficients λ may vary and cause a mismatch between the assumed distribution and the true distribution. Thereby it is beneficial to consider any possible λ for robust evaluations. We encourage more exploration of the choice of λ for the future work.

4.3. Datum Detection

Inspired from [13], by leveraging the duality theorem, we could derive the dual problem $\mathcal{W}_p(P_i, Q) := \max_{(f,g) \in C^0(\mathcal{Z})^2} \langle f, P_i \rangle + \langle g, Q \rangle$ of the primal problem in (1), where $C^0(\mathcal{Z})$ is the set of all continuous functions, $f, g \in \mathbb{R}^{m \times 1}$ are dual variables. Strong duality theorem says if π^* and (f^*, g^*) are optimal variables of the corresponding primal and dual problem respectively, then we have $\mathcal{W}_p(\pi^*) = \mathcal{W}_p(f^*, g^*)$. We applied this theorem into our scenario, when γ_i is the interpolating measure between P_i and Q , we have two separate dual problems for $\mathcal{W}_p(P_i, \gamma_i)$ and $\mathcal{W}_p(\gamma_i, Q)$, where

$$\begin{aligned} \mathcal{W}_p(P_i, Q) &= \max_{(f,g) \in C^0(\mathcal{Z})^2} \langle f, P_i \rangle + \langle g, \gamma_i \rangle \\ &+ \max_{(h,j) \in C^0(\mathcal{Z})^2} \langle h, \gamma_i \rangle + \langle j, Q \rangle. \end{aligned} \quad (9)$$

Therefore, we can get $\partial_{P_i} \mathcal{W}_p(P_i, Q) \approx \partial_{P_i} \mathcal{W}_p(P_i, \gamma_i) = (f^*)^T$, which is the gradient of the distance w.r.t. the distribution P_i . We could measure the quality of each datum using the *calibrated gradient* as follows,

$$\frac{\partial \mathcal{W}_p(P_i, \gamma_i)}{\partial P_i(z_l)} = f_l^* - \sum_{j \in \{1, \dots, m\} \setminus l} \frac{f_j^*}{m-1}, \quad (10)$$

which represents the rate of change in $\mathcal{W}_p(P_i, \gamma_i)$ w.r.t. the given datum z_l in D_i .

Interpretation This value is interpreted as the contribution of a specific datum to the distance since it determines the shifting direction based on whether it is positive or negative. If the value is positive/negative, shifting more probability mass to that datum will result in an increase/decrease of the distance between the local distribution and the interpolating measure, further resulting in an increase/decrease




Figure 2. We compare 5 different approaches for shuffled data detections: **gamma** corresponds to $\partial \mathcal{W}_p(P_i, \gamma_i)$; **eta** corresponds to $\partial \mathcal{W}_p(P_i, \eta_{Q_i})$; **access data** corresponds to $\partial \mathcal{W}_p(P_i, Q)$ (access both datasets). *Noise Portion* represents the actual noisy data ratio in a dataset, *count* (c-) indicates the number of detected negative calibrated gradient values; *Detection ratio* (r-) measures detection accuracy: ($\#$ detected noisy data / $\#$ noisy data); **servergamma** (gray) and **servereta** (purple) lines corresponds to $\partial \mathcal{W}_p(\eta_{P_i}, \eta_{Q_i})$ and $\partial \mathcal{W}_p(\eta_{P_i}, Q)$.

of the distance between the local and the target distribution. **Accurate detections** Since γ_i is the approximated interpolating measure sampling from both P_i and Q , then any function involved γ_i in (9) also provides gradient information of P_i . However, with fixed Q , the interpolating measure will shift only if P_i is changed, and in our empirical exploration in Figure 2, since the server shares η_{Q_i} to the client, $\partial \mathcal{W}_p(P_i, \eta_{Q_i})$ detects better than $\partial \mathcal{W}_p(P_i, \gamma_i)$.

4.4. Theoretical Analysis

In this section, we will provide the theoretical insights to justify our approach. The convergence guarantee and complexity analysis are similar to [28] based on the computation of Wasserstein distance in a Federated scenario.

Convergence Guarantee First we will show that Algorithm 1 has a convergence guarantee.

Theorem 1 Let P_i be the distribution of i -th client, where $i \in [1, N]$, and $Q^{(k)}$ be the Wasserstein barycenter at iteration k , $\gamma_i^{(k)}, \eta_{P_i}^{(k)}, \eta_{Q_i}^{(k)}$ be the interpolating measures computed in the Algorithm 1. Define

$$A^{(k)} = \sum_i^N \left[\mathcal{W}_p(P_i, \gamma_i^{(k)}) + \mathcal{W}_p(Q^{(k)}, \gamma_i^{(k)}) \right], \quad (11)$$

then, the sequence $(A^{(k)})$ is non-increasing and converges to $\sum_{i=1}^N \mathcal{W}_p(P_i, Q)$.

We refer the proof to Appendix 8 and also conduct toy experiments to verify it later.

Complexity Analysis FedBary algorithm computes 3 OT plans per iteration for evaluating one client, 2 for the client and 1 for the server. If there are N clients, the server totally calculates $N \times K$ OT plans. Each OT plan's complexity is based on the network simplex algorithm, which operates at $\mathcal{O}(2m^3 \log(2m))$ if balanced. Any interpolating measure between two distributions is supported by at

most $2m + 1$ points. Based on the approximation for the interpolating measure in (5), the overall computation cost is $\mathcal{O}(3NK(Sm^2 + S^2)\log(m + S))$ with the support size S for $\gamma_i^{(k)}$. In Figure 6 in the Appendix, we could observe different S will not affect the contribution topology. Thus, we could reduce the complexity with small S . In real applications, FedBary might be appropriate, especially when N is large since the complexity is linear with N . To show FedBary could be applied at scale, we compare the elapsed time for evaluating with different N and S in Table 2.

Performance Bound Without any downstream training data detections, the Wasserstein distance $\mathcal{W}_p(P_i, Q)$ can be simply used as the proxy for the validation performance when D_Q is available. If D_Q is not available on the server, we assume that D_Q can be drawn i.i.d from the Wasserstein barycenter Q approximated by Algorithm 1.

Theorem 2 ([13]) Denote $f_t : \mathcal{X} \rightarrow \{0, 1\}^V$, $f_v : \mathcal{X} \rightarrow \{0, 1\}^V$ as the labeling functions for training and validation data. where V is the number of different labels. Let $f : \mathcal{X} \rightarrow \{0, 1\}^V$ be the model trained on training data. Let $P_i(\cdot|y_o)$ and $Q(\cdot|y_o)$ be the corresponding conditional distributions given label y_o . Assume that the model f is ϵ -Lipschitz and the loss function $\mathcal{L} : \{0, 1\}^V \times [0, 1]^V \rightarrow \mathbb{R}^+$ is k -Lipschitz in both inputs. Define distance function d between (x_i, y_i) and (x_q, y_q) as (8). Under a certain cross-Lipschitzness assumption for f_i and f_v , we have

$$\begin{aligned} & \mathbb{E}_{x \sim Q(x)}[\mathcal{L}(f_v(x), f(x))] \\ & \leq \mathbb{E}_{x \sim P_i(x)}[\mathcal{L}(f_t(x), f(x))] + k\epsilon\mathcal{W}_p(P_i, Q) + \mathcal{O}(\epsilon V) \end{aligned} \quad (12)$$

The proof is outlined in Appendix 9. This theorem indicates validation loss is linearly changed with the $\mathcal{W}_p(P_i, Q)$ if the training loss (first term on rhs) is small enough.

Privacy Guarantee In our setting, privacy is guaranteed from two aspects: (1) Raw training samples cannot be reconstructed. (2) Generating data with similar appearance is not feasible. The reasons behind this are: Firstly, the stacked vector does not replicate the original data structure. Secondly, hyperparameter t for approximating the interpolating measure, the OT plan between \tilde{X}_i and γ_i , and OT plan between η_{Q_i} and η_{P_i} are only known to the client i himself. Therefore, it is challenging for attackers who need to know all the information to infer raw datasets.

5. Experiments

We demonstrate the computation of the Wasserstein barycenter among three Gaussian distributions from the empirical perspective to make Theorem 1 more convincing. For each distribution, we sample 100 data points from 2D Gaussian distributions with distinct means but the same covariance matrix. We set $t = 0.5$ for interpolating the measure, as approximated in Section 5. To assess the accuracy

N	ExactFed	GTG	MR	DataSV	Ours ($S = 100/500$)
5	31m	33s	5m	25m	1m / 2m
10	3h20m	7m	40m	2h30m	2m / 4m
50	-	-	-	-	14m / 30m
100	-	-	-	-	21m / 1h

Table 2. Evaluation time with different size of N : For ExactFed, GTG and MR, we only consider the evaluation time after model training; Evaluation time of FedBary increases linearly with N .




Figure 3. Approximated and true Wasserstein barycenter of 3 Gaussian distributions: 3-th epoch and 10-th epoch (overlapping).

of our computations, we compare the results obtained using FedBary with the barycenter approximation with data access as outlined in the work by [5]. In our study, we quantified the disparity between the barycenter estimated within FL and the barycenter derived from accessible data by computing the squared errors for each position and subsequently aggregating all differences. Remarkably, our experimental results demonstrated a swift convergence, typically within a mere 10 iterations in Figure 3.

5.1. Clients Evaluation

Datasets We used the CIFAR-10 dataset in this section for experiments (more dataset experiments are shown in Appendix), and followed the data settings in [20]. We simulate $N = 5$ clients and randomly sample data for each client with 5 cases: (1) Same Distribution and Same Size; (2) Different Distributions and Same Size; (3) Same Distribution and Different Sizes; (4) Noisy Labels and Same Size; (5) Noisy Features and Same Size. Refer Appendix 10.1 for details.

With Validation data When the validation data is available, our algorithm can be directly used to compute the distance between data from each client and the validation data, which is usually stored on the server side in reality. Thus, the shorter distance implies the better contribution from client data. We compared our method with other data valuation metrics. Baselines are *GTG-Shapley* with its variants (*GTG-Ti/GTG-Tib*) [20], *FedShapley* [41], *MR/OR* [34], *DataSV* [8] and exact calculation *exactFed*.

We visualize all cases in Figure 4 to show the percentage of contribution when the number of clients equals 5 under different valuation metrics. The X-axis is the contribution score(%), and the y-axis are valuation approaches. Each

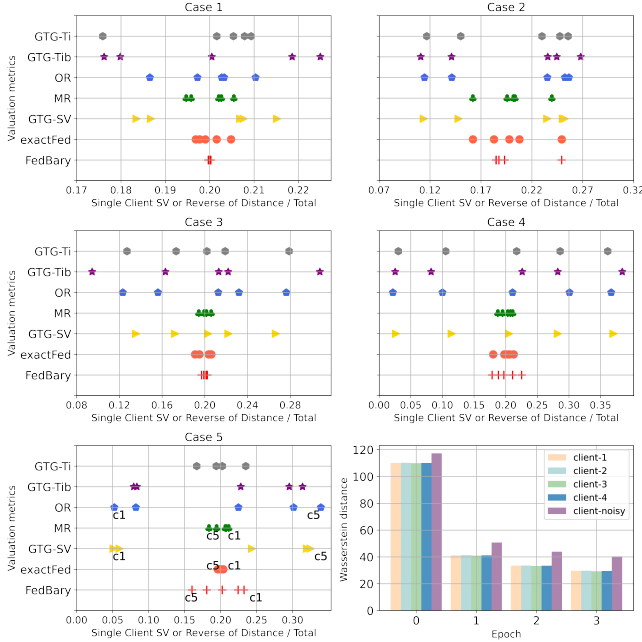


Figure 4. Scatter plots: percentage of contribution for 5 clients under different valuation metrics (Case1~5); Histogram: distance between the local distribution and the Wasserstein barycenter when validation set is not available.

marker stands for the score of a client. We measured the percentage of contribution by dividing each client’s Shapley value or inverse of distance by the total. The exactFed provides the original Shapley value, which can be considered as ground truth to some extent. For case 1, the result of our distance metric shows that the contributions for all 5 clients are almost equal, with each of them occupying around 20%, which is a signal that this metric outperforms others since the datasets are i.i.d with the same size. For case 2, FedBary could differentiate distributions and follow similar contribution scores with exactFed and MR. For case 3, with the same distribution, the size of data samples will have a trivial influence on the contribution, which shows the robustness to replications of FedBary, while other approaches GTG-Ti,GTG-Tib, GTG-SV are sensitive to the data size. We can also find MR, which is shown to be more accurate than other SV approaches for evaluations, achieves similar performance with ours and exactFed in this setting. If noise exists in labels of data (case 4), we can expect that the higher the percentage of noisy labels, the smaller the contribution score is. Some metrics like GTG-SV will show clear discretization among clients with different percentages of noisy labels. Most of them range from 2% to 38%. However, the exactFed is not that dispersed, as the SVs among different clients are close. FedBary follows this trend and delivers a result similar to exactFed and MR, showing its ability to differentiate the mislabeled data. For case 5, FedBary outperforms other approximated approaches since they will have an inverse evaluation for

the clients: with a larger proportion of noise, the contribution score is larger. FedBary is sensitive to feature noise and could capture the right ordering of client contributions: with a larger proportion of noise, the contribution score is smaller. Overall, FedBary provides better evaluations.

Without Validation data Wasserstein barycenter could assist in identifying irrelevant clients (distribution that is far from others) or distribution with noisy data points. Although there is another attempt by calculating $\mathcal{W}_p(P_i, P_j), i \neq j$ [28] to measure the data heterogeneity, such procedure needs to compute $\binom{N}{2}$ OT plans in each iteration, causing high computational cost. We simulate the scenario where there are 4 clients with i.i.d distributions and the 5-th client with noisy data, i.e., with noisy features in case (5) above. Then we approximate the Wasserstein barycenter with support $\tilde{\mathbf{X}}_Q$ among these clients and measure the distance $\mathcal{W}_p(\tilde{\mathbf{X}}_i, \tilde{\mathbf{X}}_Q)$. We plot the result in Figure 4, in which we could find that the $\mathcal{W}_p(\tilde{\mathbf{X}}_5, \tilde{\mathbf{X}}_Q)$ is larger than other distances, leading to the conclusion that this dataset is relatively irrelevant.

5.2. Noisy Feature/Mislabeled Data Detection

As aforementioned, our approach represents the first attempt in FL to perform noisy data detections without sharing any data samples. To gauge the accuracy and effectiveness of our algorithm, we conducted comparative evaluations with existing approaches that could access the dataset. Baselines here are *Lava*, *LOO* and *KNNShapley*.

We follow the experiment setting in [12], where we consider two types of synthetic noise: 1) label noise where we randomly assign the label that is different from the original label of the data points with a predefined proportion; 2) feature noise where we add standard Gaussian random noises to the original features. Here we randomly choose the proportion $p_{\text{noise}}\%$ of the training dataset to perturb. Here we consider three different levels $p_{\text{noise}} \in \{5, 15, 30\}\%$. We plot noisy feature detections ($p = 10$) with corresponding point removal experiments on CIFAR10 dataset in Figure 5. The point removal experiment is performed with the following steps: removing data points from the entire training dataset in descending order of gradient values. Following the similar example in [12], we use a logistic regression model with feature embeddings for this classification task. This process can be used to show the validation performance before and after removing training noisy points identified by FedBary. Specifically, once every more 5% of datum is removed, we fit the model with the remaining dataset and evaluate its test accuracy on the holdout dataset. We visualize the accuracy w.r.t. fraction of removed valuable datasets on the right side (b for noisy feature samples, d for mislabeled samples). As for noisy feature detection (a,b), our proposed approach is sensitive to noises in features based on the calibrated gradient, leading to supe-

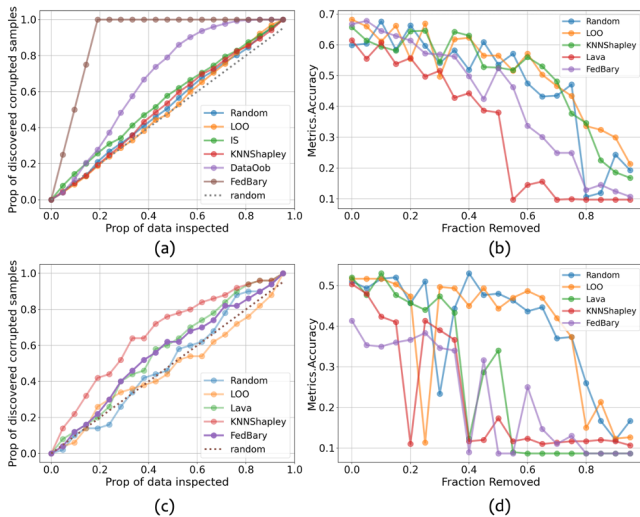


Figure 5. Detection results on CIFAR10: **Corrupt feature samples detections and point Removal Comparison** (a,b): FedBary and Lava are superior; **Misabeled samples detections and point Removal Comparison** (c,d): FedBary performs similarly to Lava and conducts relatively accurate detections;

rior performance in noisy features. The detection line of Lava is overlapping with ours thus we do not plot it (note that in Lava raw data should be accessed). We also validate our statement in (e): Notably, the number of the positive gradient in $\partial\mathcal{W}_p(P_i, \eta_Q)$ is the same as the number of noisy data, representing our approach detects all noisy data without other clean data, which convincing its effectiveness. In (b) we could find Lava could detect valuable points that affect the model performance most while FedBary is runner-up with privacy guarantee. For the label detection (c,d), we observe the Wasserstein distances change trivially with different portions of label noise, thus making the detections relatively poor. However, FedBary performs similarly with approaches that could access data, indicating our approach does not sacrifice the performance of the benchmark. In (d), Lava performs slightly better, while KNNShapley gets stuck when about 55% of the dataset is removed. LOO is relatively random in the experiment. Overall, we could say FedBary indeed provides valuable information as removing high negative values will lead to lower accuracy. It is worthy to note that, for corrupt feature detections, FedBary performs 100% success rate (only noisy data points and all of them have positive calibrated gradients) in various datasets, showing its generalization ability.

Boost FL Model We also simulate an FL setting, where 5000 training samples are divided randomly and assigned to 5 clients, in which there are 500 noisy data in total. There are 1000 validation data and 1000 testing data held by the server. Our evaluation approach proceeds as follows: before training a federated model, the server calculates the Wasserstein distance using validation data with the local samples and filters noisy samples; the filtered new training data are used to collaboratively train a federated model. We com-

Data Size	Duplication	
	Whole Dataset	One datum
5000	39.14	40.94
2×5000	39.23	40.96
3×5000	39.18	40.99

Table 3. Distance behavior under duplications

pared various datasets, and the aggregated algorithm is FedAVG [21]. Our approach could help detect noisy data and boost FL model performance in Table 4 in Appendix 11.4.

5.3. Duplication Robustness

One concern in real-world data marketplaces revolves around the ease of data duplication, which does not introduce any new information. [13, 43] have emphasized the importance of a metric that can withstand data duplication. It is also likely that the client identifies the datum with the highest contribution and duplicates it in an attempt to maximize profit. FedBary is formulated in terms of distributions and automatically disregards duplicate sets. We conducted the experiment using CIFAR10, where we simulated 5000 training data and 5000 validation data. We repeated the training set up to three times, and the distance remained unchanged. We also duplicate a single datum with a large negative gradient value in 10 multiple times for evaluation, while the result shows it would increase the distance due to the resulting imbalance in the training distribution caused by copying that particular point. We show the result in Table 3. We also discuss the time robustness in Appendix 11.2.

6. Conclusion

We propose Wasserstein distance in the FL context as a new metric for client evaluation and data detection. Compared to previous approaches, Our approach, FedBary, offers a more transparent and robust framework, substantiated through both theoretical and empirical analyses. In addition, it is designed to be applicable in real-world data marketplaces, enabling the data client evaluation before the FL training process. This facilitates the selection of only the most relevant clients and data points for the FL training process, thereby optimizing computational efficiency and enhancing model performance. We posit that such a metric holds considerable promise not only for the purpose of client evaluation but also as a cornerstone for developing incentive mechanisms within FL systems. While FedBary marks a significant advancement in this field, there remain several open questions and avenues for further research, which we have detailed in Appendix 11.5.

References

- [1] Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011. [3](#)
- [2] David Alvarez-Melis and Nicolo Fusi. Geometric dataset distances via optimal transport. *Advances in Neural Information Processing Systems*, 33:21428–21439, 2020. [2](#), [3](#), [4](#)
- [3] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005. [1](#)
- [4] Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [5] Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In *International conference on machine learning*, pages 685–693. PMLR, 2014. [3](#), [6](#)
- [6] Vitaly Feldman and Chiyuan Zhang. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891, 2020. [2](#)
- [7] R Flamary, N Courty, D Tuia, and A Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1:1–40, 2016. [2](#)
- [8] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*, pages 2242–2251. PMLR, 2019. [1](#), [2](#), [6](#), [4](#), [5](#)
- [9] Judy Hoffman, Mehryar Mohri, and Ningshan Zhang. Algorithms and theory for multiple-source adaptation. *Advances in neural information processing systems*, 31, 2018. [2](#)
- [10] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nezihe Merve Gurel, Bo Li, Ce Zhang, Costas J Spanos, and Dawn Song. Efficient task-specific data valuation for nearest neighbor algorithms. *arXiv preprint arXiv:1908.08619*, 2019. [2](#)
- [11] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. Towards efficient data valuation based on the shapley value. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1167–1176. PMLR, 2019. [1](#), [2](#)
- [12] Kevin Fu Jiang, Weixin Liang, James Zou, and Yongchan Kwon. Opendataval: a unified benchmark for data valuation. *arXiv preprint arXiv:2306.10577*, 2023. [3](#), [7](#), [4](#)
- [13] Hoang Anh Just, Feiyang Kang, Jiachen T Wang, Yi Zeng, Myeongseob Ko, Ming Jin, and Ruoxi Jia. Lava: Data valuation without pre-specified learning algorithms. *arXiv preprint arXiv:2305.00054*, 2023. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#)
- [14] Feiyang Kang, Hoang Anh Just, Anit Kumar Sahu, and Ruoxi Jia. Performance scaling via optimal transport: Enabling data selection from partially revealed sources. *arXiv preprint arXiv:2307.02460*, 2023. [2](#)
- [15] Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcev, and Gustavo K Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE signal processing magazine*, 34(4):43–59, 2017. [1](#)
- [16] Yongchan Kwon and James Zou. Beta shapley: a unified and noise-reduced data valuation framework for machine learning. *arXiv preprint arXiv:2110.14049*, 2021. [2](#)
- [17] Yongchan Kwon and James Zou. Data-oob: out-of-bag estimate as a simple and efficient data value. In *International Conference on Machine Learning*, pages 18135–18152. PMLR, 2023. [3](#)
- [18] Yongchan Kwon, Manuel A Rivas, and James Zou. Efficient computation and analysis of distributional shapley values. In *International Conference on Artificial Intelligence and Statistics*, pages 793–801. PMLR, 2021. [2](#)
- [19] Jinkun Lin, Anqi Zhang, Mathias Léculyer, Jinyang Li, Aurojit Panda, and Siddhartha Sen. Measuring the effect of training data on deep learning predictions via randomized experiments. In *International Conference on Machine Learning*, pages 13468–13504. PMLR, 2022. [2](#)
- [20] Zelei Liu, Yuanyuan Chen, Han Yu, Yang Liu, and Lizhen Cui. Gtg-shapley: Efficient and accurate participant contribution evaluation in federated learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(4):1–21, 2022. [1](#), [2](#), [6](#), [4](#), [5](#)
- [21] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. [1](#), [8](#)
- [22] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. Agnostic federated learning. In *International Conference on Machine Learning*, pages 4615–4625. PMLR, 2019. [4](#)
- [23] Eduardo Fernandes Montesuma and Fred Maurice Ngole Mboula. Wasserstein barycenter for multi-source domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16785–16793, 2021. [2](#), [4](#)
- [24] Brendan Pass. Multi-marginal optimal transport: theory and applications. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 49(6):1771–1790, 2015. [3](#)
- [25] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. [2](#), [1](#)
- [26] Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems*, 33:19920–19930, 2020. [2](#)
- [27] Alain Rakotomamonjy, Rémi Flamary, Gilles Gasso, M El Alaya, Maxime Berar, and Nicolas Courty. Optimal transport for conditional domain matching and label shift. *Machine Learning*, pages 1–20, 2022. [2](#)
- [28] Alain Rakotomamonjy, Kimia Nadjahi, and Liva Ralaivola. Federated wasserstein distance. *arXiv preprint arXiv:2310.01973*, 2023. [2](#), [3](#), [4](#), [5](#), [7](#)
- [29] Ievgen Redko, Amaury Habrard, and Marc Sebban. Theoretical analysis of domain adaptation with optimal transport. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2017, Skopje, Macedo-*

- nia, September 18–22, 2017, *Proceedings, Part II 10*, pages 737–753. Springer, 2017. 2
- [30] Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adaptation under target shift. In *The 22nd International Conference on artificial intelligence and statistics*, pages 849–858. PMLR, 2019. 2
- [31] Amirhossein Reiszadeh, Farzan Farnia, Ramtin Pedarsani, and Ali Jadbabaie. Robust federated learning: The case of affine distribution shifts. *Advances in Neural Information Processing Systems*, 33:21554–21565, 2020. 4
- [32] Rachael Hwee Ling Sim, Yehong Zhang, Mun Choon Chan, and Bryan Kian Hsiang Low. Collaborative machine learning with incentive-aware model rewards. In *International conference on machine learning*, pages 8927–8936. PMLR, 2020. 2
- [33] Rachael Hwee Ling Sim, Xinyi Xu, and Bryan Kian Hsiang Low. Data valuation in machine learning: “ingredients”, strategies, and open challenges. In *Proc. IJCAI*, pages 5607–5614, 2022. 1
- [34] Tianshu Song, Yongxin Tong, and Shuyue Wei. Profit allocation for federated learning. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 2577–2586. IEEE, 2019. 1, 6, 4, 5
- [35] Tiffany Tuor, Shiqiang Wang, Bong Jun Ko, Changchang Liu, and Kin K Leung. Overcoming noisy and irrelevant data in federated learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 5020–5027. IEEE, 2021. 2
- [36] Rosanna Turrise, Rémi Flamary, Alain Rakotomamonjy, and Massimiliano Pontil. Multi-source domain adaptation via weighted joint distributions optimal transport. In *Uncertainty in Artificial Intelligence*, pages 1970–1980. PMLR, 2022. 2
- [37] Cédric Villani. *Topics in optimal transportation*. American Mathematical Soc., 2021. 2, 3
- [38] Cédric Villani et al. *Optimal transport: old and new*. Springer, 2009. 2
- [39] Jiachen T Wang and Ruoxi Jia. Data banzhaf: A robust data valuation framework for machine learning. In *International Conference on Artificial Intelligence and Statistics*, pages 6388–6421. PMLR, 2023. 2
- [40] Shengsheng Wang, Bilin Wang, Zhe Zhang, Ali Asghar Heidari, and Huiling Chen. Class-aware sample reweighting optimal transport for multi-source domain adaptation. *Neurocomputing*, 523:213–223, 2023. 2
- [41] Tianhao Wang, Johannes Rausch, Ce Zhang, Ruoxi Jia, and Dawn Song. A principled approach to data valuation for federated learning. *Federated Learning: Privacy and Incentive*, pages 153–167, 2020. 6, 5
- [42] Xinyi Xu, Lingjuan Lyu, Xingjun Ma, Chenglin Miao, Chuan Sheng Foo, and Bryan Kian Hsiang Low. Gradient driven rewards to guarantee fairness in collaborative machine learning. *Advances in Neural Information Processing Systems*, 34:16104–16117, 2021. 2
- [43] Xinyi Xu, Zhaoxuan Wu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. Validation free and replication robust volume-based data valuation. *Advances in Neural Information Processing Systems*, 34:10837–10848, 2021. 2, 8
- [44] Jinsung Yoon, Sercan Arik, and Tomas Pfister. Data valuation using reinforcement learning. In *International Conference on Machine Learning*, pages 10842–10851. PMLR, 2020. 3