

Learn How to Query from Unlabeled Data Streams in Federated Learning

Yuchang Sun¹, Xinran Li¹, Tao Lin², Jun Zhang¹

¹Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology
²Westlake University

{yuchang.sun, xinran.li}@connect.ust.hk, lintao@westlake.edu.cn, eejzhang@ust.hk

Abstract

Federated learning (FL) enables collaborative learning among decentralized clients while safeguarding the privacy of their local data. Existing studies on FL typically assume offline labeled data available at each client when the training starts. Nevertheless, the training data in practice often arrive at clients in a streaming fashion without ground-truth labels. Given the expensive annotation cost, it is critical to identify a subset of informative samples for labeling on clients. However, selecting samples locally while accommodating the global training objective presents a challenge unique to FL. In this work, we tackle this conundrum by framing the data querying process in FL as a collaborative decentralized decision-making problem and proposing an effective solution named LeaDQ, which leverages multi-agent reinforcement learning algorithms. In particular, under the implicit guidance from global information, LeaDQ effectively learns the local policies for distributed clients and steers them towards selecting samples that can enhance the global model’s accuracy. Extensive simulations on image and text tasks show that LeaDQ advances the model performance in various FL scenarios, outperforming the benchmarking algorithms.

1 Introduction

Federated learning (FL) (McMahan et al. 2017) has emerged as a distributed paradigm for training machine learning (ML) models while preserving local data privacy. In FL, the model training process involves multiple clients, each possessing its own local training data with different distributions. These clients jointly optimize an ML model based on their local data and periodically upload the model updates to the server. Afterwards, the server updates the global model by applying the aggregated updates and distributes the current model to the clients for next-iteration training. This approach has gained significant attention in practical applications such as healthcare (Rieke et al. 2020) and Internet of Things (Zhang et al. 2022a), due to its potential to address privacy concerns associated with centralized data storage and processing.

Most of existing FL studies assume that clients have a fixed pool of training data samples with ground-truth labels (Li et al. 2023; Ye et al. 2023; Yang, Xiao, and Ji 2023;

Shi et al. 2023b), which, however, is unrealistic in many applications. In practice, the data samples usually arrive at the clients without any label (Jin et al. 2020). For example, hospitals collect raw medical images in the course of routine care which are initially unannotated raw images. As labeling these images requires highly experienced clinical staff to spend minutes per image, it is more feasible to select a subset of samples for label querying and subsequent model training. In this case, designing an effective strategy to identify the samples that could provide the most benefit in model training becomes a critical problem.

Data querying strategies for non-FL settings have been widely studied in the *active learning* (AL) literature (Settles 2009; Liu et al. 2022; Cacciarelli and Kulahci 2024). However, there are limited studies towards adapting these strategies to FL settings due to the challenges brought by the inherently decentralized nature of FL. The major challenge lies in the conflict between the collaborative goal of optimizing the global model and the limited access to local data on clients. To be specific, FL aims to optimize the global model over all the training data across clients while each client can only access its private training data. Given this fact, it is difficult for clients to select the most critical data samples which promote the global training. One straightforward approach is prioritizing local data on each client, which, however, is a myopic strategy without considering the goal of global model training, ultimately leading to suboptimal data selections. More recently, there are some attempts to explore the data querying problem in FL with fixed unlabeled dataset by designing local selection rules (Kim et al. 2023; Cao et al. 2023). Nevertheless, the performance of these methods is restricted since the data selections on clients tend to operate in a non-cooperative fashion without the guidance towards the global training goal. To solve these problems, it is imperative to design a specific data querying strategy tailored for FL with decentralized data streams while accounting for the collaborative objective of global model training.

In this work, we investigate the data querying problem in FL, particularly focusing on a case where unlabeled data arrive at clients in a streaming fashion. By alternating local data query and model training procedures, clients aim to collaboratively optimize the global model, as illustrated in Fig. 1. In the subsequent sections, we first empirically show the impact of various data querying strategies on the perfor-

mance of the global model. In particular, we find that a coordinated data querying strategy, which considers the training goal from a global perspective, yields greater performance benefit compared to individual data querying decisions. Motivated by this observation, we propose a data querying algorithm named *LeaDQ* where clients query data in a decentralized manner while considering the global objective. Specifically, we formulate the data querying problem in FL as a decentralized partially observable Markov decision process (Dec-POMDP) and leverage the multi-agent reinforcement learning (MARL) algorithms to learn local data querying policies for clients. By implicitly incorporating the status of the global model, the learned policies effectively guide the clients to make query decisions that are collaborative and ultimately beneficial for global training. We conduct simulations to compare the LeaDQ algorithm with several representative data querying strategies. The experimental results on various image and text tasks demonstrate that LeaDQ selects samples that result in more meaningful model updates, leading to improved model accuracy.¹

2 Related Works

Active learning AL algorithms are designed to select a subset of unlabeled data samples and query their labels from an oracle. The design goal is to improve the model performance on the target dataset by selecting the most informative samples within the querying cost budget (Liu et al. 2022; Cacciarelli and Kulahci 2024). The classic data selection metrics are uncertainty-based, diversity-based or the hybrid ones. Specifically, the uncertainty-based AL algorithm (Wang and Shang 2014) selects the unlabeled data with the lowest prediction confidence, which is straightforward for improving the performance of ML models. Meanwhile, the diversity-based method (Sener and Savarese 2018) aims to reduce redundancy in the selected instances for labeling, ensuring a more efficient use of labeling resources by avoiding similar samples. Furthermore, the hybrid method attempts to combine the strengths of both uncertainty and diversity in the selection criteria. In particular, some works (Fang, Li, and Cohn 2017; Zhang et al. 2022b) leverage reinforcement learning (RL) to learn a data query policy which is able to select the impactful data samples.

Federated active learning Recent studies have extended the investigations on data query problems to FL scenarios, termed as *federated active learning* (FAL) (Kim et al. 2023; Cao et al. 2023; Ahn et al. 2024; Wu et al. 2023; Kong et al. 2024). In FAL, clients select some local data samples and ask the oracles to label these samples. Kim et al. (2023) design a two-step algorithm named **LoGo** to select informative data samples while resolving the inter-class diversity on clients. Meanwhile, **KAFAL** (Cao et al. 2023) leverages the specialized knowledge by the local clients and the global model to deal with their objective mismatch. Nevertheless, the clients only have access to their own data and tend to query the samples that are beneficial for their local training objective. Because of the discrepancy between the local and

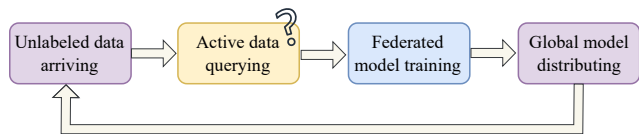


Figure 1: Workflow of the proposed LeaDQ framework.

global training objectives, the queried data may bring less benefit for the global model training. Moreover, these works assume that all unlabeled data samples are available beforehand and require evaluating them at every time. Different from these works, we focus on a more challenging stream-based setting, where the unlabeled samples sequentially arrive at clients and may be accessed only once.

In addition, another thread of works (Chen et al. 2020; Marfoq et al. 2023; Gong et al. 2023; Shi et al. 2023a) designs various approaches to deal with the unlabeled data or streaming data in FL. These works, however, are beyond the scope of this paper, since they tackle other aspects such as self-supervised learning or data storage on clients instead of the data querying problem.

3 Problem Formulation

We consider an FL system where K clients with unlabeled data streams cooperate to train an ML model by alternating data querying and model training. The goal of clients is to minimize the following objective:

$$\min_{\theta \in \mathbb{R}^M} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{P}} [\ell(\theta; \mathbf{x}, y)], \quad (1)$$

where $\ell(\theta; \mathbf{x}, y)$ denotes the loss value of model θ on data sample (\mathbf{x}, y) with data input \mathbf{x} and corresponding label y . Note that the target data distribution \mathcal{P} is unknown and inaccessible. We assume that the aggregated training data on all clients follow the same distribution as that of the target dataset. Nevertheless, the local data distribution \mathcal{P}_k of each client $k \in [K]$ is typically different from the global data distribution, which is also known as the non-independent and identically distributed (non-IID) setup in FL (Kairouz et al. 2021). Formally, the assumptions on data distributions are given as: $\cup_{k \in [K]} \mathcal{P}_k = \mathcal{P}$ and $\mathcal{P}_k \neq \mathcal{P}, \forall k \in [K]$.

Raw training data without ground-truth labels arrive at client k according to the underlying distribution \mathcal{P}_k (Marfoq et al. 2023). According to the data arrival status, the training process of optimizing model θ can be divided into R rounds. At the beginning of each round $r \in \{1, 2, \dots, R\}$, some new unlabeled data samples arrive at each client. To utilize these data for model training, clients actively query the labels of some selected samples from an oracle. Afterwards, they collaboratively train the ML model with the current labeled dataset. The above process is illustrated in Fig. 2. In the following, we introduce the procedure of data querying and model training in detail.

Active data querying In the r -th round, a set of N_u unlabeled data samples arrives at client k , denoted by $\mathcal{U}_k^r = \{\mathbf{x}_i\}_{i=1}^{N_u}$ where \mathbf{x}_i is the feature of the i -th sample. Each

¹The code is available at <https://github.com/hiyuchang/leadq/>.

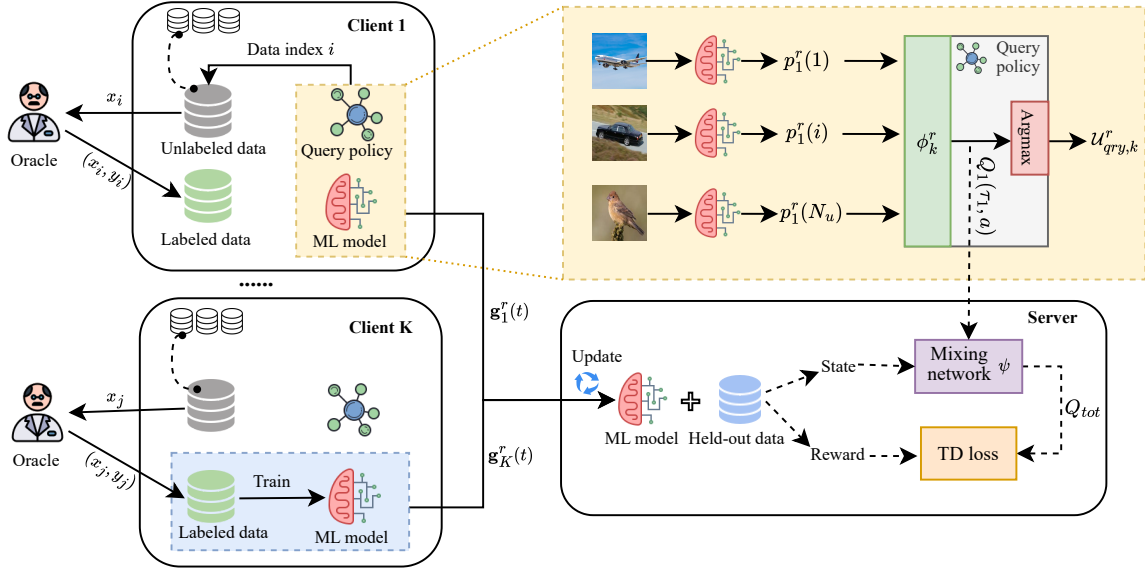


Figure 2: An overview of the proposed LeaDQ framework. The yellow block is the process of active data querying and the blue block is the process of federated model training. The dashed lines are conducted only when training the query policies.

client then makes a binary decision $\mathbf{a}_k^r = \{0, 1\}^{N_u}$, indicating whether a particular sample is selected for labeling.² Specifically, $\mathbf{a}_k^r[i] = 1$ if the i -th sample is chosen for querying and $\mathbf{a}_k^r[i] = 0$ otherwise. If being queried, an oracle provides the ground-truth label $y_i \in \mathcal{Y}$ for this sample \mathbf{x}_i . We denote the number of samples for labeling per round on each client as N_q and the queried sample set on client k as $\mathcal{U}_{qr,y,k}^r$. Afterwards, the client updates its local training set \mathcal{L}_k^{r-1} by incorporating the newly labeled data samples, i.e.,

$$\mathcal{L}_k^r = \mathcal{L}_k^{r-1} \cup \mathcal{U}_{qr,y,k}^r. \quad (2)$$

Given the high labeling cost, only a subset of unlabeled samples is queried for annotation. The labeling budget of each client is expressed as:

$$|\mathcal{U}_{qr,y,k}^r| = N_q. \quad (3)$$

The data querying policy, parameterized with a *query model* ϕ , outputs the selection decisions \mathbf{a}_k^r for the given unlabeled dataset \mathcal{U}_k^r on client k . It is worth noting that the query model may vary among clients and across rounds depending on the training stage and the current status of ML model.

Federated model training Based on the local labeled data \mathcal{L}_k^r , clients train the current global model θ^r to optimize the following objective:

$$\min_{\theta^r \in \mathbb{R}^M} f(\theta^r) \triangleq \sum_{k=1}^K \frac{|\mathcal{L}_k^r|}{|\mathcal{L}^r|} F_k(\theta^r; \mathcal{L}_k^r), \quad (4)$$

²In practical applications, the number of arrived data samples and queried samples may vary among clients or across training rounds. We assume these values are constant in this work for simplicity of expression but our method can be easily adapted to more general cases.

with the local loss function

$$F_k(\theta^r; \mathcal{L}_k^r) \triangleq \frac{1}{|\mathcal{L}_k^r|} \sum_{(\mathbf{x}, y) \in \mathcal{L}_k^r} \ell(\theta^r; \mathbf{x}, y). \quad (5)$$

Here $\mathcal{L}^r = \bigcup_{k=1}^K \mathcal{L}_k^r$ denotes the labeled dataset on all clients. As the training data contain privacy-sensitive information and cannot be exposed to the server, we adopt the classic FedAvg algorithm (McMahan et al. 2017) to optimize the ML model. Specifically, the model training process in one round can be divided into T training iterations. At the beginning, the clients initialize the current model as³ $\theta^r(0) = \theta^{r-1}$. In the t -th training iteration, client k performs local training for multiple steps to compute the accumulated gradients $\mathbf{g}_k^r(t)$. The gradients are then uploaded to the server for aggregation. The global model is then updated using these gradients as:

$$\theta^r(t+1) = \theta^r(t) - \eta \sum_{k=1}^K \mathbf{g}_k^r(t), \forall t = 0, \dots, T-1. \quad (6)$$

After T training iterations, the global model is finalized as $\theta^r = \theta^r(T)$.

Through iterative data querying and model training, the ML model is optimized by clients. However, the conflict remains between the clients' shortsighted view on its local data and the target data distribution of the global model. The queried data may be therefore redundant or not that beneficial for the global model training. To address this problem,

³Instead of using the trained model from previous round, some works also reset the model to random or designed weights for next-round training (Ash and Adams 2020).

it is intuitively helpful to provide the clients with a global view that captures the current training status and implies the desired data for facilitating the global model training. To clearly demonstrate this phenomenon, we present an example in the next section.

4 Motivating Example

In this section, we show empirically that a data querying strategy with a global view can benefit the FL training. In particular, we explore an FL setup in which clients collaborate to classify data samples from the MNIST dataset by training a LeNet model (LeCun et al. 1998). We assume there are $N_u = 10$ unlabeled data samples arriving at each client in each round and only one data sample is selected for labeling. The classical coresets method (Sener and Savarese 2018) selects the samples such that the model learned over the selected samples is competitive for the remaining data points. For an unlabeled dataset \mathcal{U} and given labeled dataset \mathcal{L} , the data sample to be queried is selected as:

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{U}} \min_{\mathbf{x}' \in \mathcal{L}} \|\mathbf{x} - \mathbf{x}'\|_2. \quad (7)$$

We first implement the coresets approach on each client locally, i.e., $\mathcal{U} = \mathcal{U}_k^r$, termed as **Local Coreset**. In Local Coreset, clients make selections based on the local available data samples. Besides, we consider an ideal approach termed as **Global Coreset** where the server can access to all clients' data $\mathcal{U} = \sum_{k \in [K]} \mathcal{U}_k^r$ and determine one queried sample for each client. After querying the labels, clients iteratively train the ML model based on the labeled dataset. The experimental details are referred to Appendix B.1.

The evaluation in Fig. 3 focuses on the effectiveness of various data querying strategies and the distribution of queried data samples. It is observed that the Global Coreset strategy improves the model accuracy compared with Local Coreset, benefiting from the utilization of global information when selecting unlabeled data. This can be verified by the distribution divergence between the queried samples in each round and the global data distribution as shown in Fig. 3 (b). In particular, Global Coreset can select critical samples that coincide with the target data distribution, since it leverages the information from a global view.

Nevertheless, the global information in Global Coreset requires computing the sample-wise distance between each unlabeled sample and labeled sample. As the data samples are distributed on different clients, it is impractical to directly implement the Global Coreset, which severely violates the data privacy of clients. Thus, it becomes crucial to devise a strategy that leverages global insights without necessitating direct access to data across all clients. To this end, we propose the LeaDQ algorithm under a centralized training decentralized execution (CTDE) paradigm, which implicitly incorporates the global objective into the querying strategies while maintaining the decentralized execution as in FL. Consequently, LeaDQ achieves comparable performance (as shown in Fig. 3) to Global Coreset while preserving the data privacy.




Figure 3: Comparisons between different data querying strategies. Left: The model accuracy is computed on the test data. Right: The *distribution divergence* is computed as the Kullback–Leibler (KL) divergence between the distributions of selected unlabeled data on all clients and the target data.

5 LeaDQ: Learning the Data Querying Policies for Clients

Effective data querying strategies in FL should exhibit the ability to select unlabeled samples that are beneficial for global model training, while making the decisions in a decentralized fashion. Motivated by the capabilities of MARL algorithms in making individual decisions while optimizing a collaborative objective (Vinyals et al. 2017; Wang et al. 2021), we propose an MARL-based data querying algorithm named **LeaDQ**, short for *Learn Data Querying*. In the following, we first formulate the data querying problem in FL as a Markov decision process and elaborate on the details of system design in Section 5.1. Subsequently, we detail the process of decentralized data querying and centralized policy training in Sections 5.2 and 5.3, respectively.

5.1 Data Querying in FL as A Decentralized Decision-Making Process

The objective of data querying per client k in FL aims to select samples from local unlabeled dataset \mathcal{U}_k^r to query their labels in round r . By data querying and model training, these clients aim to collaboratively optimize the global model θ^r . However, each client only has access to its local data without the knowledge of either data distribution of other clients or the target global data distribution, i.e., the decision making process is based on partial observation in its nature.

To describe such a setup, we frame the data querying problem as a decentralized partially observable Markov decision process (Dec-POMDP) (Oliehoek and Amato 2016) denoted by a tuple $\langle \mathcal{S}, \mathcal{A}, P, R, \mathcal{O}, K, \gamma \rangle$. We view each querying round as a discrete timestep and each client as an agent in this process.

At each discrete timestep r , a **global state** $s^r \in \mathcal{S}$ reveals the current status of global model training. We assume the server owns a held-out dataset $\mathcal{D}_{\text{held}}$ (Fang and Ye 2022; Huang, Ye, and Du 2022) and the global state is defined as the *prediction confidence* of the global model θ^r on it:

$$s^r = \left[\max_{y \in \mathcal{Y}} p(y|\mathbf{x}; \theta^r) \right]_{\mathbf{x} \in \mathcal{D}_{\text{held}}}. \quad (8)$$

We assume the held-out dataset follows the same distribution as that of the whole dataset, and thus the global state

defined in (8) serves as a low-cost indicator to show the current model's performance on the target dataset and can effectively reveal the current training status.

When a set of unlabeled data \mathcal{U}_k^r with the length of N_u arrives, each client k gets the **local observation** $\mathbf{o}_k^r \in \mathcal{O}$ of these data. Inspired by previous works (Wang and Shang 2014), we use the *predictive logits* output by the current model θ^r as the local observation, i.e.,

$$\mathbf{o}_k^r = [l(\mathbf{x}; \theta^r)]_{\mathbf{x} \in \mathcal{U}_k^r}, \quad (9)$$

where $l(\mathbf{x}; \theta^r) = p(y|\mathbf{x}; \theta^r)$, $\forall y \in \mathcal{Y}$ encapsulates the predicted logits. The local observation \mathbf{o}_k^r is a low-dimension representation of arrived samples in \mathcal{U}_k^r and also reflects the uncertainty of samples. It is intuitive that samples with higher levels of uncertainty would be more helpful to model updates when being included in the training dataset.

Based on its local observation \mathbf{o}_k^r , each client chooses **actions** $\mathbf{a}_k^r = \{0, 1\}^{N_u}$ to select samples for labeling. Here the binary decision $\mathbf{a}_k^r[i] = 1$ denotes the case of the i -th sample being chosen for querying and $\mathbf{a}_k^r[i] = 0$ otherwise. The joint action is then given by $\mathbf{A}^r = [\mathbf{a}_1^r, \mathbf{a}_2^r, \dots, \mathbf{a}_K^r]$. After querying unlabeled data, the clients cooperatively update the ML model based on the labeled dataset following the procedure elaborated in Section 3. Such evolution of the ML model's states can be characterized by the transition function $P(s'|s, \mathbf{a}) : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$.

Afterwards, the updated global model returns to clients a joint **reward** R^r , which reveals the impact of selected actions on the ML model training. As the collaborative goal of clients is to improve the overall performance of the ML model θ^r , this reward should reveal its current training status. Thus we define the reward as the difference of model accuracy on the held-out data before and after updating, i.e.,

$$R^r \equiv R(\mathbf{s}^r, \mathbf{A}^r) = \text{Acc}(\theta^r; \mathcal{D}_{\text{held}}) - \text{Acc}(\theta^{r-1}; \mathcal{D}_{\text{held}}). \quad (10)$$

The reward in (10) shapes a *global* view of the model training, which helps clients to make better *local* decisions of data querying. Formally, clients aim to find the optimal data querying policies that maximize the total discounted reward $R_{tot}^r = \sum_{j=r}^{r+J} \gamma^j R^j$ with episode length J and discount γ .

By trial-and-error, the data query policies are directed to incorporate the information of the global model training, indicated by state \mathbf{s}^r and reward R^r , into the collaborative decisions, ultimately promoting the global model training. When executing the query policies, it is sufficient for clients to select local unlabeled data samples in a decentralized manner. In summary, such CTDE strategy bridges the gap between the local data querying and the objective of global model training, as detailed in the following subsections.

5.2 Decentralized Policy Execution for Querying Data Samples

In round r , client k has access to its local action-observation history $\tau_k \in \mathcal{T} \equiv (\mathcal{O} \times \mathcal{A})^*$. Based on such history, the client makes a decision of selecting unlabeled data following its data query policy. Specifically, each client feeds the local observation \mathbf{o}_k^r of arrived unlabeled data into its policy

Algorithm 1: The LeadQ framework

```

1: for each round  $r = 1$  to  $R$  do
2:   // Active data querying
3:   for each client  $k = 1$  to  $K$  do
4:     Compute predictive logits as local observations  $\mathbf{o}_k^r$ 
       {▷ Eq. 9}
5:     Compute local Q-values as  $Q_k(\mathbf{a}, \mathbf{o}_k^r; \phi_k^r)$ 
6:     Choose actions  $\mathbf{a}_k^r$  according to greedy policy
       {▷ Eq. 11}
7:     Query the oracle for label  $y_i$  of data sample  $\mathbf{x}_i$  with
        $a_k^r[i] = 1$ 
8:     Update the labeled dataset as  $\mathcal{L}_k^r$            {▷ Eq. 2}
9:   end for
10:  // Federated model training
11:  Initialize the global model as  $\theta^r(0) = \theta^{r-1}$ 
12:  for each iteration  $t = 0$  to  $T - 1$  do
13:    for each client  $k = 1$  to  $K$  do
14:      Train the current model  $\theta^r(t)$  based on the local
        labeled dataset  $\mathcal{L}_k^r$ 
15:      Upload the gradients  $\mathbf{g}_k^r(t)$  to the server
16:    end for
17:    The server updates the global model as  $\theta^r(t + 1)$ 
        {▷ Eq. 6}
18:  end for
19:  Update the global model as  $\theta^r = \theta^r(T)$ 
20: end for

```

network, denoted by ϕ_k^r , to output the local Q-value. Subsequently, it chooses the action that induces the highest local Q-value, which is given by:

$$\mathbf{a}_k^r = \arg \max_{\mathbf{a}: |\mathbf{a}| = N_q} Q_k(\mathbf{a}, \mathbf{o}_k^r; \phi_k^r), \quad (11)$$

where $|\mathbf{a}| = N_q$ implies N_q unlabeled samples are queried. Depending on the value of \mathbf{a}_k^r , each unlabeled sample is determined for querying or not. As such, the clients use the query policy to select data in a decentralized manner without violating the data privacy. The details of the proposed framework are summarized in Algorithm 1.

When querying data, the clients rely on their local policies, which shall guide them to select critical samples for the global model training. To integrate the global information into this process, we adapt the QMIX algorithm (Rashid et al. 2018) into the LeadQ framework for training policies.

5.3 Centralized Policy Training for Learning Local Policies

By selecting unlabeled data, all clients contribute to the global model training and get a shared reward, serving as the global feedback for the joint decision of clients. However, it is difficult for clients to directly use the shared reward for local data selections. To link the global feedback and local decisions, we define the joint Q-value as a weighted sum of the local Q-values, which is given by (Rashid et al. 2018):

$$Q_{tot} = \psi(Q_1, \dots, Q_K, \mathbf{s}^r). \quad (12)$$

Here ψ is a *mixing network* from a family of monotonic functions to model the complex interactions between collaborating clients. This mixing network helps to coordinate




Figure 4: Model accuracy (%) in different data querying rounds.

the clients’ local policies while accounting for the state of current ML model \mathbf{s}^r .

We use a replay buffer \mathcal{B} to store the transition tuple $\langle \mathbf{s}, \mathbf{o}, \mathbf{a}, R, \mathbf{s}', \mathbf{o}' \rangle$, where the state \mathbf{s}' is observed after taking the action \mathbf{a} in state \mathbf{s} and receiving reward R . During policy training, we iteratively sample batches of transitions from the replay buffer and update the policy network by minimizing the temporal difference (TD) loss:

$$\begin{aligned} \mathcal{L}(\{\phi_k\}, \psi) &= \mathbb{E}_{(\mathbf{s}, \mathbf{o}, \mathbf{a}, R, \mathbf{s}', \mathbf{o}') \sim \mathcal{B}} \left[\left(y^{tot} - Q_{tot}(\mathbf{s}, \mathbf{o}, \mathbf{a}; \{\phi_k\}, \psi) \right)^2 \right], \end{aligned} \quad (13)$$

where $y^{tot} = R + \gamma \max_{\mathbf{u}} Q_{tot}(\mathbf{s}', \mathbf{o}', \mathbf{u}; \{\phi_k\}, \psi)$ with discount γ . Indeed, minimizing the TD loss in (13) optimizes the local policies $\{\phi_k\}$ on clients towards a global goal, i.e., increasing the reward R . As such, by following their learned policies, clients can select the beneficial unlabeled data for global model training. The pseudo code of policy training is deferred to Appendix A.1.

6 Experiments

6.1 Setup

We simulate an FL system with one server and $K = 10$ clients. We evaluate the algorithms on two image classification tasks, i.e., SVHN (Netzer et al. 2011) and CIFAR-100 (Krizhevsky and Hinton 2009), and one text classification task, i.e., 20Newsgroup (Lang 1995). We train a convolutional neural network (CNN) model with four convolutional layers (Kim et al. 2023) on the SVHN dataset, a ResNet-18 (He et al. 2016) on the CIFAR-100 dataset, and a DistilBERT (Sanh et al. 2020) model on the 20Newsgroup dataset. To simulate the non-IID setting, we allocate the training data to clients according to the Dirichlet distribution with concentration parameter $\alpha = 0.5$ (Li et al. 2022). In each round, $N_u = 10$ unlabeled data samples arrive at each client independently and each client selects $N_q = 1$ data sample for label querying. The details of datasets are summarized in Table 1. More implementation details can be

Type	Dataset	Non-IID Type
Text	20Newsgroup	Distribution-based
Image	SVHN	Distribution-based Quantity-based
	CIFAR-100	Distribution-based

Table 1: Summary of datasets.

found in Appendix B.2. Extended results can be found in Appendix C. In all tables, the best performances are highlighted in **bold**, and the second-best ones are underlined.

Baselines We compare the proposed algorithm with the following baselines. First, we adapt several representative AL methods to FL, including: (i) **Uncertainty** (Wang and Shang 2014): Each client selects the samples with lowest prediction confidence for label querying; (ii) **Coreset** (Sener and Savarese 2018): Each client selects the samples such that the model learned over these samples is competitive for the remaining unlabeled data points. Second, two state-of-the-art federated AL algorithms are also compared, including: (iii) **LoGo** (Kim et al. 2023): Each client selects samples using a clustering-based sampling strategy; (iv) **KAFAL** (Cao et al. 2023): Each client selects samples with the highest discrepancy between the specialized knowledge by the local model on clients and the global model.

6.2 Main Results

We show the model accuracy in different querying rounds on several datasets in Fig. 4. We observe that the proposed LeaDQ algorithm achieves the best accuracy in all the querying rounds. On the SVHN dataset, the naive local querying strategies (i.e., Uncertainty and Coreset) attain weaker performance than other approaches, showcasing the importance of designing specific strategies for FL. On the CIFAR-100 dataset, the baselines have similar performance, while LeaDQ surpasses the best-performance baseline at the later




Figure 5: Visualization of the data samples in SVHN dataset ($R = 100$). Left: Label distribution of local data on a random client, target global data, and queried data by LeadDQ. Right: T-SNE plots (Van der Maaten and Hinton 2008) of feature distributions of the global data and queried data by LeadDQ.

N_q/N_u	1/30	1/20	1/10	2/10	3/10
Uncertainty	44.35	50.64	59.28	73.31	79.30
Coreset	39.96	45.07	62.72	75.97	79.66
LoGo	45.45	49.94	59.66	76.02	79.86
KAFAL	41.88	49.30	<u>64.67</u>	76.32	<u>80.30</u>
LeadDQ	45.68	52.31	65.02	<u>76.28</u>	80.37

Table 2: Model accuracy (%) on SVHN with different values of arrived samples N_u and queried samples N_q .

training stage by around 6%, implying its advantages in querying critical samples for global model training.

Visualization of queried data samples To visualize the querying strategies provided by LeadDQ, we show the queried samples distributions in Fig. 5. According to Fig. 5 (left), the queried unlabeled data samples have similar distributions with the target global data distribution, despite the fact that the local data follows a different distribution with the global data. Besides, the t-SNE plots in Fig. 5 (right) also demonstrate that LeadDQ is able to select unlabeled data that have a similar feature distribution with the target global data.

6.3 Results With Diverse Scenarios

Results with various arrived and queried data samples

We show the model accuracy with different numbers of arrived samples N_u and queried samples N_q per round in Table 2. We observe from the results that as the ratio of queried data samples N_q/N_u increases, the model achieves higher accuracy with all the data querying strategies. This is because leveraging more labeled data samples on clients effectively improves the model’s performance. In most cases, LeadDQ outperforms the baselines due to its advantages in identifying critical samples. Meanwhile, when the number of queried samples is larger ($N_q > 1$), LeadDQ and KAFAL have competitive performances, surpassing other strategies.

Results with different setups of data heterogeneity

In addition, we simulate two types of data heterogeneity, i.e., the distribution-based skew and quantity skew (Li et al. 2022). The results are illustrated in Tables 3 and 4, respec-

	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1.0$
Uncertainty	48.09	59.28	63.54
Coreset	49.77	62.72	65.04
LoGo	43.70	59.66	59.97
KAFAL	<u>52.14</u>	<u>64.67</u>	<u>66.48</u>
LeadDQ	53.98	65.02	66.50

Table 3: Model accuracy (%) on SVHN with different degrees of non-IID under distribution-based skew. α is the value of concentration parameter in Dirichlet distribution, and a larger α indicates that the data distribution across clients is closer to IID.

	$C = 2$	$C = 3$	$C = 4$
Uncertainty	62.26	67.73	67.10
Coreset	<u>64.15</u>	<u>67.77</u>	<u>68.43</u>
LoGo	59.69	63.15	67.07
KAFAL	20.19	20.32	21.75
LeadDQ	65.34	68.19	69.06

Table 4: Model accuracy (%) on SVHN with different degrees of non-IID under quantity skew. C is the number of classes on each client, and a larger C indicates that the data distributions across clients are closer to IID.

tively, in which we show the model accuracy after $R = 500$ rounds when applying different data querying strategies. We vary the data heterogeneity among clients by adjusting the concentration parameter α and the number of classes on each client C . When the data tend to be more heterogeneous, the model performance degrades as the training becomes harder. We find that LeadDQ achieves higher performance compared with the baselines in all cases, showing its effectiveness in finding beneficial samples for global model training. However, the FAL baselines may even fail in identifying informative samples and thus lead to suboptimal performance. Overall, these results demonstrate the robustness and adaptability of the proposed LeadDQ algorithm.

7 Conclusions

In this paper, we investigate the data querying problem in stream-based federated learning, highlighting the conflict between local data access and global querying objective. We introduce a novel MARL-based algorithm named LeadDQ, which learns local data query policies on distributed clients. LeadDQ promotes cooperative objectives among clients, ultimately leading to improved machine learning model training. Extensive experiments validate the superiority of LeadDQ over state-of-the-art baselines, demonstrating its effectiveness in querying critical data samples and enhancing the quality of the ML model. We note that LeadDQ provides a framework for data querying in FL. Future research could explore the incorporation of more advanced algorithms beyond FedAvg and QMIX to enhance performance further.

Acknowledgements

This work was supported by the Hong Kong Research Grants Council under the Areas of Excellence scheme grant AoE/E-601/22-R and NSFC/RGC Collaborative Research Scheme grant CRS_HKUST603/22.

References

- Ahn, J.; Ma, Y.; Park, S.; and You, C. 2024. Federated Active Learning (F-AL): An Efficient Annotation Strategy for Federated Learning. *IEEE Access*, 12: 39261–39269.
- Ash, J. T.; and Adams, R. P. 2020. On Warm-Starting Neural Network Training. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- Cacciarelli, D.; and Kulahci, M. 2024. Active learning for data streams: A survey. *Mach. Learn.*, 113(1): 185–239.
- Cao, Y.; Shi, Y.; Yu, B.; Wang, J.; and Tao, D. 2023. Knowledge-Aware Federated Active Learning with Non-IID Data. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 22222–22232.
- Chen, Y.; Ning, Y.; Slawski, M.; and Rangwala, H. 2020. Asynchronous Online Federated Learning for Edge Devices with Non-IID Data. In *2020 IEEE International Conference on Big Data (BigData)*, 15–24.
- Fang, M.; Li, Y.; and Cohn, T. 2017. Learning how to Active Learn: A Deep Reinforcement Learning Approach. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Fang, X.; and Ye, M. 2022. Robust Federated Learning with Noisy and Heterogeneous Clients. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10062–10071.
- Gong, C.; Zheng, Z.; Wu, F.; Shao, Y.; Li, B.; and Chen, G. 2023. To Store or Not? Online Data Selection for Federated Learning with Limited Storage. In Ding, Y.; Tang, J.; Sequeda, J. F.; Aroyo, L.; Castillo, C.; and Houben, G., eds., *Proceedings of the ACM Web Conference 2023 (WWW)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Huang, W.; Ye, M.; and Du, B. 2022. Learn from Others and Be Yourself in Heterogeneous Federated Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10133–10143.
- Jin, Y.; Wei, X.; Liu, Y.; and Yang, Q. 2020. Towards utilizing unlabeled data in federated learning: A survey and prospective. arXiv:2002.11545.
- Kairouz, P.; et al. 2021. Advances and open problems in federated learning. *Found. Trends Mach. Learn.*, 14(1–2): 1–210.
- Kim, S.; Bae, S.; Song, H.; and Yun, S. 2023. Re-Thinking Federated Active Learning Based on Inter-Class Diversity. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kong, X.; Zhang, W.; Qu, Y.; Yao, X.; and Shen, G. 2024. FedAWR: An Interactive Federated Active Learning Framework for Air Writing Recognition. *IEEE Trans. Mob. Comput.*, 23(5): 6423–6436.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. [Online]. Available: <https://www.cs.toronto.edu/~kriz/cifar.html>.
- Lang, K. 1995. NewsWeeder: Learning to Filter Netnews. In *Proceedings of the 12th International Conference on Machine Learning (ICML)*, 331–339.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11): 2278–2324.
- Li, Q.; Diao, Y.; Chen, Q.; and He, B. 2022. Federated Learning on Non-IID Data Silos: An Experimental Study. In *38th IEEE International Conference on Data Engineering (ICDE)*, 965–978.
- Li, Z.; Lin, T.; Shang, X.; and Wu, C. 2023. Revisiting weighted aggregation in federated learning with neural networks. In *International Conference on Machine Learning*, 19767–19788. PMLR.
- Liu, P.; Wang, L.; Ranjan, R.; He, G.; and Zhao, L. 2022. A Survey on Active Deep Learning: From Model Driven to Data Driven. *ACM Comput. Surv.*, 54(10s): 221:1–221:34.
- Marfoq, O.; Neglia, G.; Kameni, L.; and Vidal, R. 2023. Federated Learning for Data Streams. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 8889–8924.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 7. Granada, Spain.
- Oliehoek, F. A.; and Amato, C. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer Briefs in Intelligent Systems.
- Rashid, T.; Samvelyan, M.; de Witt, C. S.; Farquhar, G.; Foerster, J. N.; and Whiteson, S. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, 4292–4301.
- Rieke, N.; Hancox, J.; Li, W.; Milletari, F.; Roth, H. R.; Albarqouni, S.; Bakas, S.; Galtier, M. N.; Landman, B. A.; Maier-Hein, K. H.; Ourselin, S.; Sheller, M. J.; Summers, R. M.; Trask, A.; Xu, D.; Baust, M.; and Cardoso, M. J. 2020. The future of digital health with federated learning. *npj Digit. Medicine*, 3.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv:1910.01108.

Sener, O.; and Savarese, S. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *6th International Conference on Learning Representations (ICLR)*.

Settles, B. 2009. Active learning literature survey.

Shi, J.; Wu, Y.; Zeng, D.; Tao, J.; Hu, J.; and Shi, Y. 2023a. Self-Supervised On-Device Federated Learning From Un-labeled Streams. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 42(12): 4871–4882.

Shi, Y.; Liang, J.; Zhang, W.; Tan, V. Y. F.; and Bai, S. 2023b. Towards Understanding and Mitigating Dimensional Collapse in Heterogeneous Federated Learning. In *The Eleventh International Conference on Learning Representations (ICLR)*.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).

Vinyals, O.; Ewalds, T.; Bartunov, S.; Georgiev, P.; Vezhnevets, A. S.; Yeo, M.; Makhzani, A.; Küttler, H.; Agapiou, J.; Schrittwieser, J.; et al. 2017. Starcraft II: A new challenge for reinforcement learning. arXiv:1708.04782.

Wang, D.; and Shang, Y. 2014. A new active labeling method for deep learning. In *2014 International Joint Conference on Neural Networks (IJCNN)*, 112–119.

Wang, J.; Xu, W.; Gu, Y.; Song, W.; and Green, T. C. 2021. Multi-agent reinforcement learning for active voltage control on power distribution networks. *Advances in Neural Information Processing Systems*, 34: 3271–3284.

Wu, X.; Pei, J.; Chen, C.; Zhu, Y.; Wang, J.; Qian, Q.; Zhang, J.; Sun, Q.; and Guo, Y. 2023. Federated Active Learning for Multicenter Collaborative Disease Diagnosis. *IEEE Trans. Medical Imaging*, 42(7): 2068–2080.

Yang, Y.; Xiao, P.; and Ji, K. 2023. SimFBO: Towards Simple, Flexible and Communication-efficient Federated Bilevel Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Ye, R.; Xu, M.; Wang, J.; Xu, C.; Chen, S.; and Wang, Y. 2023. FedDisco: Federated Learning with Discrepancy-Aware Collaboration. In *International Conference on Machine Learning (ICML)*.

Zhang, T.; Gao, L.; He, C.; Zhang, M.; Krishnamachari, B.; and Avestimehr, A. S. 2022a. Federated Learning for the Internet of Things: Applications, Challenges, and Opportunities. *IEEE Internet Things Mag.*, 5(1): 24–29.

Zhang, Y.; Tong, H.; Xia, Y.; Zhu, Y.; Chi, Y.; and Ying, L. 2022b. Batch Active Learning with Graph Neural Networks via Multi-Agent Deep Reinforcement Learning. In *Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*.