KNOWING WHEN TO ADOPT THE SHARED KNOWLEDGE: BENEFIT-CALIBRATED FEDERATED LORA UNDER PARTIAL PARTICIPATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Federated low-rank adaptation (FedLoRA) enables communication-efficient fine-tuning without centralizing client data. However, under partial participation, heterogeneous data, and resource constraints, federated collaboration may not benefit all clients simultaneously. Moreover, aggregation may conceal negative transfer, while short-term local signals do not align with long-term benefits reliably. We therefore revisit FedLoRA from a client-benefit perspective and investigate when knowledge should be shared, personalized, or deferred. We define federation benefit relative to local proximal training and propose a benefit-calibrated framework that accumulates evidence across communication rounds and make the decision by Bayesian method. In particular, historical participation is tracked to prevent persistent client under-representation, while each client's estimated benefit, associated uncertainty, and resource budget jointly determine whether collaboration follows a shared, personalized, or deferred decision. The analysis across multiple datasets shows that the benefits of federation vary substantially across client, motivating client-level evaluation and evidence-based collaboration. Experimental results demonstrate that our method achieves stronger and more consistent performance than existing baselines in most of scenarios.

1 INTRODUCTION

Adapting pre-trained models to downstream applications is often constrained by the computation, memory, and communication resources available at the training site (Han et al., 2024). Low-rank adaptation (LoRA) alleviates these constraints by representing model updates with compact low-rank factors, thereby reducing the number of trainable parameters and the cost of model adaptation (Hu et al., 2022). LoRA typically assumes that the data required for adaptation are accessible within a single training environment. However, in many practical scenarios, relevant data are distributed across organizations or edge devices and cannot be directly centralized because of privacy, regulatory, or ownership constraints (Arbaoui et al., 2024). Federated learning provides a natural mechanism for collaborative training to solve such problem by keeping raw data locally and exchanging only model-related information (McMahan et al., 2017). Integrating LoRA into this collaborative process gives rise to federated LoRA (FedLoRA), which supports parameter-efficient adaptation over decentralized data while reducing client-side training and communication costs (Yang et al., 2025).

However, real-world federated systems are generally characterized by partial client participation, non-identically distributed data, heterogeneous computation and communication resources (Ye et al., 2023). Under these constraints, enabling collaborative adaptation does not ensure that every participating client will benefit from the resulting federation. Existing FedLoRA studies have primarily sought to improve global model accuracy or convergence by designing more effective aggregation strategies, accommodating heterogeneous adaptation capacities, or mitigating discrepancies among local updates (Wu et al., 2024; Yi et al., 2026). Although these approaches improve collaborative training from a system-wise perspective, they potentially conceal substantial differences among clients. In particular, knowledge sharing may improve some clients while causing negative transfer for others whose local distributions or objectives are poorly aligned with the aggregated model (Tan et al., 2023; Zhang et al., 2023). Moreover, short-term clients' losses and update similarities may not

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

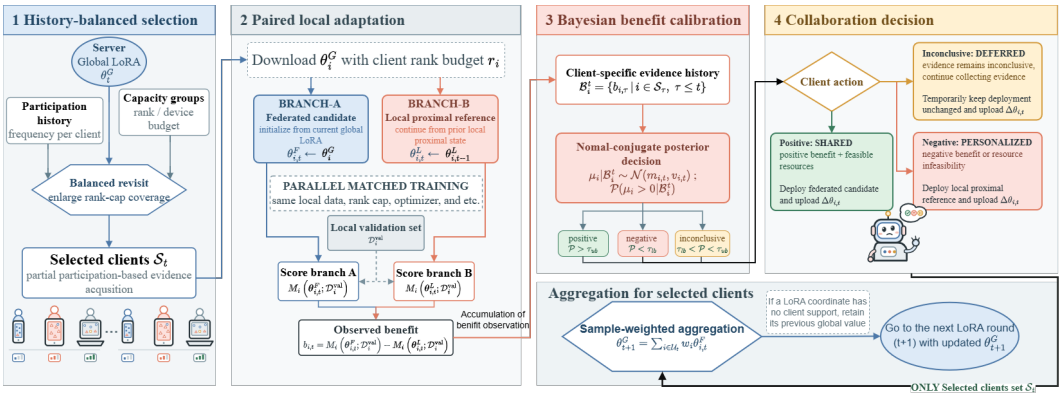

Figure 1: Overview of BC-FedLoRA. Each selected client trains a federated candidate and a local reference under matched conditions. Their validation-performance difference is accumulated across participation rounds to estimate the client's federation benefit and its uncertainty. Positive evidence leads to sharing, negative evidence leads to personalization, and insufficient evidence defers deployment changes.

reliably predict long-term collaboration benefit, especially when observations are sparse since the constrained resources leading to partial participation of federated learning. A uniform collaboration policy may therefore share information with clients that are unlikely to benefit, while overlooking uncertainty in the estimated benefit and differences in client resource budgets. These limitations motivate our critical question, i.e., which clients benefit from federation, and how to make the decision for knowledge of current round?

To answer the question, we propose benefit-calibrated FedLoRA (BC-FedLoRA), a Bayesian-embedded framework that organizes federated adaptation around the expected benefit of collaboration for each client, and the overview of the framework is shown in Figure 1. Instead of applying a uniform sharing policy for all clients, BC-FedLoRA uses local proximal adaptation as a reference and accumulates evidence about the client's benefit of federation across communication rounds. Historical participation is tracked to prevent persistently underrepresented clients from receiving unreliable benefit estimates. Subsequently, we process the observed benefit values using a Bayesian posterior model, thereby incorporating the uncertainty into following decision. Based on the client's resource budget, our framework then determines whether collaborative knowledge of clients should be shared to aggregation, retained by personalized adaptation, or deferred until more evidence is available. Our design allows collaboration decisions to evolve as additional observations are collected, aiming to limit clients' exposure to potentially harmful sharing. To validate the effectiveness of our proposed framework, we evaluate BC-FedLoRA on image and text datasets under different simulation settings, comparing it with multiple baselines, including local-only training, standard FedLoRA, heterogeneous-capacity aggregation, and personalized federated learning (Sun et al., 2023; Liu et al., 2026; Chen et al., 2026). In addition to aggregate predictive performance, our evaluation measures the proportion of clients that benefit from federation, the incidence of negative transfer, tail-client performance, and communication efficiency, providing a more complete assessment of both system-level effectiveness and client-level outcomes.

Our contributions are listed as follows:

- We formulate benefit-aware collaboration in heterogeneous FedLoRA as an online decision problem under partial observations, using a continuously trained local proximal branch as a reference in paired adaptation to observe the clients' benefit and make the decision for personalization.
- We develop BC-FedLoRA, which adopts the history-balanced participation of clients, accumulates evidence across communication rounds, and uses Bayesian benefit estimates to decide whether each client should follow shared adaptation, personalized adaptation, or deferred collaboration.
- We establish a client-level evaluation protocol across multiple datasets, measuring not only aggregate performance but also negative transfer rates, tail-client outcomes, decision calibration, and

communication cost. Our ablation experiments further examine when commonly used short-term evidence and update-based signals fail to predict long-term collaboration benefit.

2 Related Works

2.1 Federated Parameter-Efficient Fine-Tuning

LoRA reduces fine-tuning cost by freezing pretrained weights and learning compact low-rank updates (Hu et al., 2022). This property has motivated a growing interest on federated parameter-efficient fine-tuning (PEFT), where communication and client-side optimization must remain feasible without centralizing private data. SLoRA improves federated PEFT under heterogeneous data through data-driven initialization (Babakniya et al., 2023), while FFA-LoRA freezes one factor to reduce communication and stabilize aggregation (Sun et al., 2024). Moreover, resource heterogeneity has been addressed by allowing client-specific adapter ranks. HetLoRA combines rank self-pruning with sparsity-aware aggregation (Cho et al., 2024), and FlexLoRA aggregates effective updates before using singular value decomposition to redistribute adapters at different ranks (Bai et al., 2024). A complementary branch of research studies the algebraic mismatch introduced by separately averaging LoRA factors. FLoRA represents the weighted sum of client updates through factor stacking (Wang et al., 2024), FedEx-LoRA compensates for the residual between factor-wise and update-wise aggregation (Singhal et al., 2025), and LoRA-FAIR jointly refines aggregation and cross-round initialization (Bian et al., 2025). Other methods retain client-specific adaptation through dual adapters or selective sharing, including FDLoRA (Lu et al., 2026), FedSA-LoRA (Guo et al., 2025), and FedALoRA (Yi et al., 2025). These advances substantially improve the efficiency, aggregation fidelity, and personalization of federated LoRA. Their primary objective, however, is typically to improve global or average personalized performance under a prescribed collaboration policy. Particularly, existing researches do not consider each client's long-term benefit and the uncertainty of benefit estimation to decide what collaboration scheme should be adopted.

2.2 Client-Level Benefit and Selective Collaboration

Personalized federated learning mitigates statistical heterogeneity by replacing a single uniform model with client-dependent models or collaboration patterns. For example, Ditto improves the personalized models by training a proximity with the help of shared global model (Li et al., 2021), FedAMP strengthens interaction between similar clients (Huang et al., 2021), and FedFomo learns client-specific model combinations from their estimated local utility (Zhang et al., 2021). Besides, there are selective approaches determine which clients or parameters should interact. FedCAC combines parameter sensitivity with client similarity (Wu et al., 2023b), FEDORA learns whether and how parameters should propagate between clients (Wu et al., 2023a), and FedSC constructs a collaboration group for each client before hierarchically fusing shared and local knowledge (Xu et al., 2025). Moreover, negative transfer has also been investigated at aggregation. ConFREE removes conflicting update components (Zheng et al., 2025), while merit-based aggregation assigns larger weights to collaborators that better match a target objective (Tupitsa et al., 2026). Existing research indicates that comprehensive and unified collaboration scheme can be less effective than personalized solutions (Allouah et al., 2024), and client-level analyses further reveal that federated averages can obscure substantial heterogeneity, with some clients deriving little benefit, or even experiencing performance degradation relative to local training (Corbucci et al., 2025). Recent trust-aware collaboration further uses example agreement and confidence to determine when public consensus is useful (Abourayya et al., 2025). Nevertheless, instantaneous validation gains, similarity, update compatibility, and predictive confidence may not identify a client's long-term benefit from federation, particularly when participation is intermittent. Our research therefore characterizes each client's collaboration benefit relative to local proximal adaptation as an uncertain observation across communication rounds, accumulates these observations into the evidence of whether collaboration is beneficial, assesses the reliability of evidence by Bayesian posterior, and uses it to choose among shared adaptation, personalization, and deferral rather than imposing collaboration whenever a client participates.

3 METHOD

We consider a federated learning system includes a central server and N clients indexed by $i \in \mathcal{N}$. Client i holds a private dataset $\mathcal{D}_i$ consisting of local training set $\mathcal{D}_i^{\text{tr}}$, validation set $\mathcal{D}_i^{\text{val}}$, and test set $\mathcal{D}_i^{\text{te}}$. The distribution of client data is different from each other, i.e., $\mathcal{P}_i \neq \mathcal{P}_j$ for $i \neq j$. Due to the constraints on computing and communication capacity, the systems typically employ a partial participation strategy for training. Specifically, the server selects only a subset $\mathcal{S}_t \subseteq \mathcal{N}$ of clients for fine-tuning in communication round t, and we adapt a pretrained model using LoRA (Hu et al., 2022) in the following. For a target weight matrix $\mathbf{W}_\ell^0$ at LoRA layer ℓ, the adaptation is given by

$$\mathbf{W}_\ell = \mathbf{W}_\ell^0 + s_\ell \mathbf{B}_\ell \mathbf{A}_\ell, \tag{1}$$

where $\mathbf{W}_\ell^0$ is the frozen component, $\mathbf{A}_\ell$ and $\mathbf{B}_\ell$ are trainable low-rank components, and s_ℓ is a scaling factor. Moreover, each client has a resource-dependent rank budget r_i, and $\boldsymbol{\theta}_t^G$ denotes the global LoRA parameters maintained by the server at the round t.

The objective of our research is to determine how individual clients should handle the knowledge from server and adaptively reduce the negative transfer. The proposed BC-FedLoRA framework consists of history-balanced client selection, paired local adaptation, Bayesian benefit estimation, and collaboration decision.

3.1 HISTORY-AWARE CLIENT PARTICIPATION

Since the limited system resource constrains the scale of participating clients, the partial participation is generally adopted in practical federated system, which means that client's federation benefit can only be observed when a it is selected to join the adaptation. Therefore, repeatedly selecting the same clients results in two potential problems. Firstly, the clients with limited participation will accumulate insufficient evidence for reliable benefit estimation. Secondly, the resulting collaboration decisions may primarily reflect frequently available clients and overlook clients from different resource groups. To mitigate these problems, BC-FedLoRA uses historical participation information to improve the coverage of benefit observations. This mechanism helps to determine where additional evidence is highly required.

Specifically, let $h_{i,t}$ ($h_{i,1} = 0$) denote the number of times that client i has been selected before round t, and N_{sup} refers to the number of clients that can be supported by the current round, which is determined by the system budget. We also assume that the resource capacity of client i is g_i, which represents the LoRA capacity that its device can support. For a candidate client set $\mathcal{S}$, where $|\mathcal{S}| = K$, we define its coverage as

$$C(\mathcal{S}) = |\{g_i : i \in \mathcal{S}\}|. \tag{2}$$

Since a larger $C(\mathcal{S})$ allows the server to acquire benefit observations from a broader range of resource conditions, the core of participant selection is to maximize the coverage. Moreover, to prevent a small number of clients from dominating the historical evidence, we minimize $\max_{i\in\mathcal{S}} h_{i,t}$ to discourage the selection of a subset containing highly overrepresented clients and minimize $\sum_{i\in\mathcal{S}} h_{i,t}$ to favor the clients with fewer historical observations. Since the label distribution of clients also affects system coverage, BC-FedLoRA uses distributional divergence as another criterion for coverage maximization. We define

$$\overline{D}_{\text{JS}}(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{i\in\mathcal{S}} D_{\text{JS}}\left(\boldsymbol{\pi}_i \,\|\, \boldsymbol{\pi}\right), \tag{3}$$

where D_{JS} is the Jensen-Shannon divergence, $\boldsymbol{\pi}_i$ denotes the empirical distribution of client i, and $\boldsymbol{\pi}$ refers to the reference distribution over the federated population. This criterion provides a deterministic preference among the subsets to provide a larger coverage. Based on the above discussion, the selected client set can be obtained through the following lexicographic optimization problem

$$\mathcal{S}_t \in \underset{\mathcal{S}\subseteq\mathcal{N},\,|\mathcal{S}|=K}{\operatorname{argmin}} \left(-C(\mathcal{S}),\ \max_{i\in\mathcal{S}} h_{i,t},\ \sum_{i\in\mathcal{S}} h_{i,t},\ \overline{D}_{\text{JS}}(\mathcal{S}) \right). \tag{4}$$

The lexicographic optimization means that BC-FedLoRA first minimizes the resource-group coverage, then minimizes the largest historical participation followed by the total participation, and finally the label-distribution divergence. After client selection, the participation history is updated as

$$h_{i,t+1} = h_{i,t} + \mathbb{I}[i \in \mathcal{S}_t], \tag{5}$$

where $\mathbb{I}[\cdot]$ is the indicator function.

The purpose of this mechanism is to distribute benefit observations more evenly across clients and resource conditions, thereby providing sufficient temporal evidence for subsequent collaboration decisions.

3.2 PAIRED LOCAL ADAPTATION

To provide reliable evidence for collaboration decisions, we compare the deployable alternatives available to each client within the same communication round. Existing works use the local-only model as a reference for quantifying client's benefit (Yu et al., 2022; Corbucci et al., 2025). However, using such method results in the local branch to drift excessively when trained on a small amount of skewed local data. To mitigate this model drift while still allowing the personalized model to adapt to local data, the local training objective is added a regularization toward the same global model. Therefore, under the same training setup, BC-FedLoRA constructs two paired branches of adaptation for each selected client, i.e., the federated branch and local proximal branch.

To help the following introduction, we use $\boldsymbol{\theta}_t^G$ to denote the global LoRA parameters received at round t, and $\boldsymbol{\theta}_{i,t-1}^L$ is the personal state retained by client i from its previous participation. We assume both branches use the client's resource-dependent LoRA rank budget. The federated branch is initialized from the current global parameters and locally optimized according to

$$\boldsymbol{\theta}_{i,t}^F \leftarrow \arg\min_{\boldsymbol{\theta}} \mathcal{L}_i\left(\boldsymbol{\theta}, \mathcal{D}_i^{\mathrm{tr}}; r_i\right), \quad \boldsymbol{\theta} \text{ initialized at } \boldsymbol{\theta}_t^G. \tag{6}$$

This branch represents the update that client i would contribute to the federated model.

In parallel, the local proximal branch starts at the client's previously retained personal state $\boldsymbol{\theta}_{i,t-1}^L$, and its local objective combines the task loss with a proximal penalty centered at the current global LoRA parameters, which is given by

$$\boldsymbol{\theta}_{i,t}^L \leftarrow \arg\min_{\boldsymbol{\theta}} \left[\mathcal{L}_i\left(\boldsymbol{\theta}, \mathcal{D}_i^{\mathrm{tr}}; r_i\right) + \frac{\lambda}{2}\left\|\boldsymbol{\theta} - \boldsymbol{\theta}_t^G\right\|_2^2\right], \quad \boldsymbol{\theta} \text{ initialized at } \boldsymbol{\theta}_{i,t-1}^L, \tag{7}$$

where $\lambda \geq 0$ controls the strength of global regularization. This initialization preserves client-specific knowledge across participation rounds, while the proximal term discourages the personal model from drifting excessively from the evolving global model(Li et al., 2021). In the following, we use the federated branch and local proximal branch to serves as the round-matched comparator for the estimation of federation benefit. Based on the same validation set of client, we evaluate the performance of local updates as $\mathcal{M}_i^{\mathrm{val}}\left(\boldsymbol{\theta}_{i,t}^F; \mathcal{D}_i^{\mathrm{val}}\right)$ and $\mathcal{M}_i^{\mathrm{val}}\left(\boldsymbol{\theta}_{i,t}^L; \mathcal{D}_i^{\mathrm{val}}\right)$, where $\mathcal{M}_i^{\mathrm{val}}$ denotes the task-specific validation metric, such as classification accuracy. The observed benefit of choosing the shared candidate over proximal personalization is defined as

$$\hat{b}_{i,t} = \mathcal{M}_i^{\mathrm{val}}\left(\boldsymbol{\theta}_{i,t}^F; \mathcal{D}_i^{\mathrm{val}}\right) - \mathcal{M}_i^{\mathrm{val}}\left(\boldsymbol{\theta}_{i,t}^L; \mathcal{D}_i^{\mathrm{val}}\right). \tag{8}$$

A positive $b_{i,t}$ favors the federated candidate for current participation, whereas a negative value favors the local proximal candidate.

3.3 BAYESIAN BENEFIT ESTIMATION AND COLLABORATION DECISION

The benefit observation $b_{i,t}$ in equation 8 describes the local training outcome at one communication round. However, the single observation may be affected by stochastic local optimization, validation noise, and the evolving global model. Therefore, making a collaboration decision directly from the latest observation results in unstable switching of adaptation. To solve this problem, our BC-FedLoRA treats each observation as a noisy measurement of a latent federation benefit and updating its posterior distribution when the client participates in adaptation, with details given in the following.

First of all, we denote the benefit observations accumulated by client i up to round t as

$$\hat{\mathcal{B}}_i(t) = \left\{\hat{b}_{i,t} \,\middle|\, i \in \mathcal{S}_t,\ t \in \mathcal{T}_i(t)\right\}, \tag{9}$$

where $\mathcal{T}_i(t)$ is the set that records client's participation. If the client is not selected at round t, no observation is introduced and the posterior remains unchanged. Furthermore, we model the benefit as

$$b_{i,t} \mid \beta_i \sim \mathcal{N}\left(\beta_i, \sigma_b^2\right), \quad t \in \mathcal{T}_i(t), \tag{10}$$

where β_i represents the latent advantage of deploying the federated branch rather than the local proximal branch for client i, and σ_b^2 refers to the observation noise. The Gaussian prior for describing the latent benefit is formulated as

$$\beta_i \sim \mathcal{N}\left(\mu_0, \tau_0^2\right). \tag{11}$$

We use a neutral prior centered at $\mu_0 = 0$ to express that no initial preference between two adaptation branches. The prior variance τ_0^2 controls the initial uncertainty, while σ_b^2 determines how strongly each new observation influences the estimate. Since the Gaussian prior is conjugate to the likelihood in equation 10, the posterior also has a Gaussian form, i.e.,

$$\beta_i \mid \hat{\mathcal{B}}_i(t) \sim \mathcal{N}\left(\mu_{i,t}, \tau_{i,t}^2\right), \tag{12}$$

where the variance and mean are respectively given by

$$\tau_{i,t}^2 = \left(\frac{1}{\tau_0^2} + \frac{|\hat{\mathcal{B}}_i(t)|}{\sigma_b^2}\right)^{-1}, \tag{13}$$

$$\mu_{i,t} = \tau_{i,t}^2 \left(\frac{\mu_0}{\tau_0^2} + \frac{1}{\sigma_b^2} \sum_{t \in \mathcal{T}_i(t)} \hat{b}_{i,t}\right). \tag{14}$$

The posterior estimation combines the prior knowledge with all paired observations collected for the client. The posterior variance decreases as additional observations are accumulated, allowing the framework to distinguish a weak estimate based on limited participation from a consistent pattern supported by repeated evidence. Based on the above model, BC-FedLoRA computes the posterior probabilities of positive and negative federation benefit.

$$p_{i,t}^{+} = \Pr\left(\beta_i > 0 \mid \hat{\mathcal{B}}_i(t)\right) = \Phi\left(\frac{\mu_{i,t}}{\tau_{i,t}}\right), \tag{15}$$

$$p_{i,t}^{-} = \Pr\left(\beta_i < 0 \mid \hat{\mathcal{B}}_i(t)\right) = \Phi\left(-\frac{\mu_{i,t}}{\tau_{i,t}}\right), \tag{16}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. A large $p_{i,t}^{+}$ indicates strong posterior evidence that the federated branch is preferable, whereas a large $p_{i,t}^{-}$ indicates evidence in favor of local proximal branch. When neither probability is sufficiently large, the observations do not support a reliable conclusion. Given a posterior-confidence threshold τ_{thr}, the decision induced by benefit estimation can be summarized as

$$a_{i,t}\left(\boldsymbol{\theta}_{i,t}\right) = \begin{cases} \textsc{Share}\ (\boldsymbol{\theta}_{i,t}^{F}), & p_{i,t}^{+} \geq \tau_{\text{thr}}, \\ \textsc{Personalize}\ (\boldsymbol{\theta}_{i,t}^{L}), & p_{i,t}^{-} \geq \tau_{\text{thr}}, \\ \textsc{Defer}\ (\boldsymbol{\theta}_{i,t-1}), & \text{otherwise}, \end{cases} \tag{17}$$

where $\boldsymbol{\theta}_{i,t}$ represents the deployed model of client i at round t. The SHARE action deploys the federated branch while PERSONALIZE action uses the independently trained local proximal branch. The DEFER action temporarily retains the previously deployed model without immediately switching to new branches, and it continues to collect the evidence and makes a decision after the confidence threshold is reached.

To enable the global model to continuously acquire new knowledge from the heterogeneous data of clients, all adaptation participants upload their federated branches after making decisions. The client $i \in \mathcal{S}_t$ transmits the LoRA update $\boldsymbol{\Delta}_{i,t}^{F} = \boldsymbol{\theta}_{i,t}^{F} - \boldsymbol{\theta}_{t}^{G}$ to server. Since the heterogeneous devices generally support different LoRA ranks, this heterogeneity of rank coordinate motivates a rank-masked FedAvg. In particular, we use $z_{i,q} \in \{0, 1\}$ to indicate whether parameter coordinate q is active under client i's rank budget. Let $\phi_i = |\mathcal{D}_i^{\text{tr}}|$ be the number of local training samples. The server aggregates LoRA updates as

$$\overline{\Delta}_{t,q} = \begin{cases} \dfrac{\sum_{i \in \mathcal{S}_t} \phi_i z_{i,q} \Delta_{i,t,q}^{F}}{\sum_{i \in \mathcal{S}_t} \phi_i z_{i,q}}, & \displaystyle\sum_{i \in \mathcal{S}_t} \phi_i z_{i,q} > 0, \\ 0, & \text{otherwise}. \end{cases} \tag{18}$$

The global LoRA parameter at the coordinate p is then updated by

$$\theta_{t+1,q}^{G} = \theta_{t,q}^{G} + \overline{\Delta}_{t,q}. \tag{19}$$

The overall algorithm is summarized in Appendix A.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Datasets, models, and baselines. We conduct experiments on datasets covering both vision and text domains. The vision benchmarks include CIFAR-100 (Krizhevsky, 2009) and Tiny-ImageNet-200 (Le & Yang, 2015), while the language benchmarks include AG News (Zhang et al., 2015), and IMDB (Maas et al., 2011). We use an ImageNet-pretrained Swin-T for the vision datasets (Liu et al., 2021), and use DistilBERT for the text datasets (Sanh et al., 2019). The pretrained backbone is frozen, while the LoRA parameters and task-specific classification head are optimized. We compare BC-FedLoRA with Local-only training, FedAvg-LoRA (Sun et al., 2023), FedProx-LoRA (Li et al., 2020), FlexLoRA (Bai et al., 2024), FedEx-LoRA (Singhal et al., 2025), Ditto-LoRA (Li et al., 2021), and FedFomo-LoRA (Zhang et al., 2021). All methods use matched data partitions, client participation schedules, LoRA capacity budget, sample limits, and random seeds 72–76.

Federated protocol and implementation details. Unless otherwise stated, each experiment contains 12 clients, from which 4 clients participate in each communication round. We construct non-IID client partitions using Dirichlet label allocation with concentration $\alpha \in \{0.1, 0.5\}$, where a smaller α represents stronger statistical heterogeneity. To model resource heterogeneity, we consider two different capacity settings with mean LoRA rank budgets of 1.5 and 2.5, respectively. Client-specific ranks follow the profiled resource plans and are bounded by a maximum rank of 8. All methods are trained for 40 communication rounds with one local epoch for each selected client and local batch size of 8. We cap the local training and test sets at 192 samples per client and reserve 20% of each capped training set for client-specific validation. The LoRA scaling factor, learning rate, and weight decay are set to 16, 10^{-3}, and 10^{-4}, respectively. We use the proximal coefficient of $\mu = 0.1$ and a four-round warm-up period. The Gaussian benefit prior has mean $\mu_0 = 0$ and standard deviation $\tau_0 = 1$, the observation standard deviation is set as $\sigma_b = 1$, and the posterior decision threshold is $\tau_{\text{thr}} = 0.85$. The experiments are executed on NVIDIA GeForce RTX 4090 GPUs.

4.2 MAIN RESULTS

In this section, we compare the mean client accuracy to evaluate the performance of the proposed framework. Moreover, other experimental results including negative transfer rate and the client decisions in BC-FedLoRA will also be provided to enhance the understanding of our proposed method. The results of client accuracy are demonstrated by the mean and standard deviation over multiple experiments. The experiments of each dataset are based on four different settings regarding to data heterogeneity and rank budget, i.e., $(\alpha = 0.1, \bar{r} = 1.5)$, $(\alpha = 0.1, \bar{r} = 2.5)$, $(\alpha = 0.5, \bar{r} = 1.5)$, and $(\alpha = 0.5, \bar{r} = 2.5)$, where $\bar{r}$ is the mean of client rank budget.Bold and underlined entries denote the best and second-best result in each comparison, respectively. We provide one representative result for each modality, and the remaining results are provided in Appendix B.1.

4.2.1 EVALUATION ON VISION DATASETS

Table 1: Mean client accuracy (%) on CIFAR-100.

Method	$(\alpha = 0.1, \bar{r} = 1.5)$	$(\alpha = 0.1, \bar{r} = 2.5)$	$(\alpha = 0.5, \bar{r} = 1.5)$	$(\alpha = 0.5, \bar{r} = 2.5)$
Local-only	74.44 ± 0.52	73.90 ± 0.49	54.54 ± 1.86	55.15 ± 2.17
FedAvg-LoRA	67.39 ± 1.17	66.57 ± 1.03	71.10 ± 0.69	71.34 ± 0.59
FedProx-LoRA	67.52 ± 1.01	66.64 ± 1.41	71.01 ± 0.87	71.47 ± 0.42
FlexLoRA	65.20 ± 1.16	65.19 ± 1.04	70.55 ± 0.96	71.44 ± 0.49
FedEx-LoRA	65.31 ± 1.30	65.30 ± 1.04	70.86 ± 0.59	71.47 ± 0.72
Ditto-LoRA	$\underline{80.43 \pm 0.72}$	$\underline{80.56 \pm 0.37}$	$\underline{71.37 \pm 1.40}$	$\underline{72.36 \pm 2.20}$
FedFomo-LoRA	74.57 ± 0.68	74.39 ± 1.11	54.78 ± 1.84	54.63 ± 1.62
BC-FedLoRA	$\mathbf{81.22 \pm 0.72}$	$\mathbf{81.35 \pm 0.70}$	$\mathbf{74.82 \pm 1.57}$	$\mathbf{74.39 \pm 1.81}$

Table 1 and Figure 2 shows the results obtained on CIFAR-100. Compared with other baselines, our BC-FedLoRA converges at a high rate and to a better result. Our method has best performance in

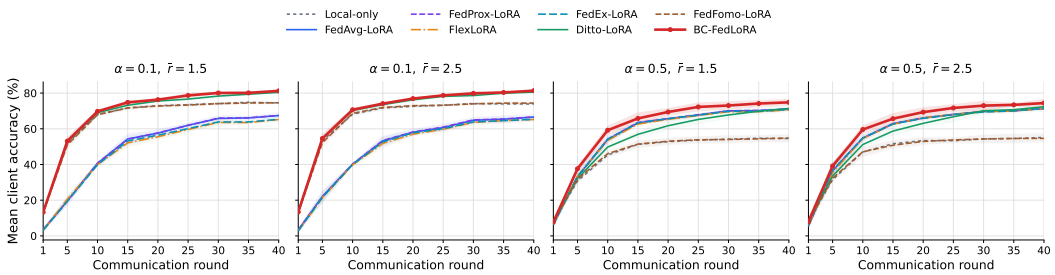

Figure 2: The convergence performance of all methods on CIFAR-100.

all settings and exceeds the strongest baseline by 0.79%, 0.79%, 3.45%, and 2.03%, respectively. The results shown in Figure 2 validate the excellent convergence performance of BC-FedLoRA. The above results on vision dataset show that adapting to different collaboration regimes rather than relying on a uniformly shared or personalized policy would benefit the federated learning.

4.2.2 Evaluation on text datasets

As for the experiments on AG News, BC-FedLoRA leads in three of four settings by 0.31%, 0.29%, and 0.40%, as shown in the Table 2. Figure 3 demonstrates the convergence performance, and it is obviously that our proposed method still converges at a high rate and to a better result.

Table 2: Mean client test accuracy (%) on AG News.

Method	$(\alpha = 0.1, \bar{r} = 1.5)$	$(\alpha = 0.1, \bar{r} = 2.5)$	$(\alpha = 0.5, \bar{r} = 1.5)$	$(\alpha = 0.5, \bar{r} = 2.5)$
Local-only	95.65 ± 1.37	95.37 ± 1.09	90.53 ± 1.26	90.54 ± 0.78
FedAvg-LoRA	87.92 ± 0.99	88.56 ± 1.63	88.24 ± 1.10	88.11 ± 1.28
FedProx-LoRA	87.68 ± 1.14	89.23 ± 1.76	88.39 ± 1.02	88.21 ± 1.36
FlexLoRA	86.94 ± 1.35	87.73 ± 1.24	88.57 ± 1.60	88.11 ± 1.23
FedEx-LoRA	86.51 ± 0.69	86.88 ± 2.03	87.64 ± 2.72	87.71 ± 1.35
Ditto-LoRA	$\underline{95.72 \pm 1.49}$	$\underline{95.73 \pm 1.48}$	$\underline{92.09 \pm 0.70}$	$\mathbf{92.02 \pm 0.47}$
FedFomo-LoRA	95.49 ± 1.45	95.64 ± 1.29	91.45 ± 0.74	91.28 ± 0.71
BC-FedLoRA	$\mathbf{96.03 \pm 1.10}$	$\mathbf{96.02 \pm 0.90}$	$\mathbf{92.49 \pm 0.45}$	$\underline{91.95 \pm 0.73}$

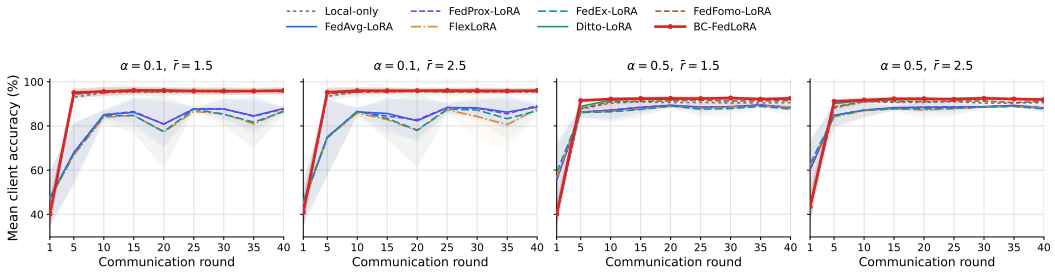

Figure 3: The convergence performance of all methods on AG News.

Across all experiments on multiple datasets, BC-FedLoRA has the highest mean client accuracy in 14 of the 16 experiments. The above results validate the performance advantages and convergence stability of our method. In the following, we investigate the tail performance and negative transfer of the proposed framework.

4.3 Tail Performance and Negative Transfer

For each dataset, we use baseline with the highest mean accuracy as the reference to relatively measure the 10th-percentile client accuracy of BC-FedLoRA. Table 3 reports accuracy differences

Table 3: Client's accuracy differences from the best baseline.

Dataset	Metric	$\alpha = 0.1$, $\bar{r} = 1.5$	$\alpha = 0.1$, $\bar{r} = 2.5$	$\alpha = 0.5$, $\bar{r} = 1.5$	$\alpha = 0.5$, $\bar{r} = 2.5$
CIFAR-100	ΔP10	$+1.30$	$+0.17$	$+3.13$	$+2.18$
	NTR reduction	0.00	$+1.67$	0.00	0.00
Tiny-ImageNet-200	ΔP10	$+0.40$	-0.06	$+0.46$	-0.44
	NTR reduction	0.00	$+1.67$	0.00	0.00
AG News	ΔP10	$+0.01$	$+0.54$	$+1.01$	$+1.06$
	NTR reduction	$+11.67$	-1.67	$+8.33$	-3.33
IMDB	ΔP10	$+1.76$	$+0.04$	$+0.75$	-2.68
	NTR reduction	$+3.33$	$+8.33$	-6.67	$+6.67$

between our BC-FedLoRA and this reference. Moreover, we also define the negative-transfer rate (NTR) relative to the matched Local-only result, and the table also reports the difference between $\mathrm{NTR}_{\mathrm{ref}}$ and $\mathrm{NTR}_{\mathrm{ours}}$ to demonstrate the corresponding characteristics. A positive value of the entry means we favor the BC-FedLoRA. BC-FedLoRA improves the 10th-percentile accuracy in 13 of 16 experiments, but the NTR result is more sophisticated. BC-FedLoRA reduces NTR in 7 experiments, leaves it unchanged in 6 experiments, and increases it in 3 experiments. Experimental results demonstrate that, in most of scenarios, BC-FedLoRA can mitigate the issue of negative transfer at the client.

4.4 Ablation Studies

We conduct ablation studies to identify which parts of BC-FedLoRA dominates its behavior. The study is based on CIFAR-100 and AG News under $\alpha = 0.1$. Every ablation retains the same client partitions, participant sequence, and resource budget as the main experiments. Table 4 reports the change in client's mean accuracy relative to BC-FedLoRA, a negative value indicates that removing or replacing the corresponding mechanism is detrimental. The full ablation discussion is provided in Appendix B.3.

Table 4: Change of client's mean accuracy (%).

Dataset	Budget	HistInt	Latest	RandSel	RandDep	GateUp	LocalRef
CIFAR-100	$\bar{r} = 1.5$	-0.56	-0.14	-0.38	-0.13	-5.28	-0.31
	$\bar{r} = 2.5$	-0.01	-0.21	-0.30	-0.14	-3.80	-0.72
AG News	$\bar{r} = 1.5$	-0.12	-0.05	-0.09	-0.16	-0.64	-0.02
	$\bar{r} = 2.5$	-0.24	-0.10	-0.14	-0.18	-0.87	-0.13

5 Conclusion

This work revisits federated LoRA from the perspective of individual client benefit rather than assuming that collaboration is always helpful. We introduced BC-FedLoRA, which estimates each client's benefit relative to a continuously trained proximal branch, accumulates evidence under partial participation, and uses posterior uncertainty to choose among shared, personalized, and deferred adaptation. Across the evaluated vision and text datasets, BC-FedLoRA achieves stronger and more consistent client-level performance than the baselines in most experiments, frequently improving mean and lower-tail accuracy while reducing negative transfer to a certain extent. The ablations further show that history-balanced selection, benefit-induced deployment and continued upload of selected clients' shared branches are important to the model behavior. These findings support treating federated adaptation as an evidence-based, client-dependent decision process.

AI USE STATEMENT

In this work, we used generative AI tools for writing and correcting experimental code, organizing experimental results, and polishing the language of the manuscript.

We have not used generative AI tools for the selection of research topic, the design of the proposed method, and the selection and review of references.

We have reviewed all AI-assisted code, analyses, and text.

We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. All authors have read and committed to upholding the policy throughout the submission and review process. Our research presents a federated LoRA framework that evaluates and calibrates the benefit of collaboration for individual clients without centralizing their raw data. All experiments use open-source benchmark datasets intended for research, and no new data involving human participants were collected.

Nevertheless, federated learning does not eliminate privacy risks, since exchanged model updates may still expose information about local data, and BC-FedLoRA does not provide a formal differential-privacy or security guarantee. We therefore report client-level and lower-tail performance, negative-transfer rates, and decision uncertainty in addition to average accuracy, while using participation history to reduce persistent underrepresentation. In practice, BC-FedLoRA should be accompanied by application-specific privacy protections, bias audits, and human oversight in sensitive settings.

We believe that improving the transparency and selectivity of federated collaboration offers meaningful benefits, while the remaining risks are identifiable and can be mitigated through appropriate safeguards.

REPRODUCIBILITY STATEMENT

We describe the algorithm, datasets, preprocessing, client partitions, resource settings, training configuration, and evaluation metrics in the paper. Experiments use paired partitions and participation schedules when they are required for comparison, and result files are checked against the frozen protocol and source manifests.

REFERENCES

Amr Abourayya, Jens Kleesiek, Bharat Rao, and Michael Kamp. Whom to trust? adaptive collaboration in personalized federated learning, 2025. URL https://arxiv.org/abs/2507.00259.

Youssef Allouah, Abdellah El Mrini, Rachid Guerraoui, Nirupam Gupta, and Rafael Pinot. Fine-tuning personalization in federated learning to mitigate adversarial clients. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 100816–100844. Curran Associates, Inc., 2024. doi: 10.52202/079017-3197. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/b6d0df730c5976ad918bbf4fb30afe7d-Paper-Conference.pdf.

Meriem Arbaoui, Mohamed-el-Amine Brahmia, Abdellatif Rahmoun, and Mourad Zghal. Federated learning survey: A multi-level taxonomy of aggregation techniques, experimental insights, and future frontiers. *ACM Trans. Intell. Syst. Technol.*, 15(6), November 2024. ISSN 2157-6904. doi: 10.1145/3678182. URL https://doi.org/10.1145/3678182.

Sara Babakniya, Ahmed Elkordy, Yahya Ezzeldin, Qingfeng Liu, Kee-Bong Song, MOSTAFA EL-Khamy, and Salman Avestimehr. SLoRA: Federated parameter efficient fine-tuning of language models. In *International Workshop on Federated Learning in the Age of Foundation Models*

in Conjunction with NeurIPS 2023, 2023. URL https://openreview.net/forum?id=06quMTmtRV.

Jiamu Bai, Daoyuan Chen, Bingchen Qian, Liuyi Yao, and Yaliang Li. Federated fine-tuning of large language models under heterogeneous tasks and client resources. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 14457–14483. Curran Associates, Inc., 2024. doi: 10.52202/079017-0461. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/1a134b50202088aa8c595cc99b310e5a-Paper-Conference.pdf.

Jieming Bian, Lei Wang, Letian Zhang, and Jie Xu. LoRA-FAIR: federated LoRA fine-tuning with aggregation and initialization refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3737–3746, October 2025.

Zeyu Chen, Yun Ji, Peng He, Bowen Wang, Liwen Shi, Zijie Zeng, and Sheng Zhang. Per-fedlora: Personalized federated lora fine-tuning for multi-task large language models. In *2026 International Wireless Communications and Mobile Computing (IWCMC)*, pp. 1801–1806, 2026. doi: 10.1109/IWCMC69287.2026.11580063.

Yae Jee Cho, Luyang Liu, Zheng Xu, Aldi Fahrezi, and Gauri Joshi. Heterogeneous LoRA for federated fine-tuning of on-device foundation models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 12903–12913, Miami, Florida, USA, nov 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.717. URL https://aclanthology.org/2024.emnlp-main.717/.

Luca Corbucci, Xenia Heilmann, and Mattia Cerrato. Benefits of the federation? analyzing the impact of fair federated learning at the client level. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, pp. 2232–2248, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400714825. doi: 10.1145/3715275.3732152. URL https://doi.org/10.1145/3715275.3732152.

Pengxin Guo, Shuang Zeng, Yanran Wang, Huijie Fan, Feifei Wang, and Liangqiong Qu. Selective aggregation for low-rank adaptation in federated learning. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 99003–99027, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/f53a37f820d5be5930415d964f4a0187-Paper-Conference.pdf.

Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning for large models: A comprehensive survey. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=lIsCS8b6zj.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized cross-silo federated learning on Non-IID data, 2021. URL https://arxiv.org/abs/2007.03797.

Alex Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 05 2009.

Y. Le and X. Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 2015.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks, 2020. URL https://arxiv.org/abs/1812.06127.

Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine*

Learning Research, pp. 6357–6368. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/li21h.html.

Muqing Liu, Chongjie Si, and Yuheng Jia. FlexloRA: Entropy-guided flexible low-rank adaptation. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=tqnkbdYWWm.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9992–10002, 2021. doi: 10.1109/ICCV48922.2021.00986.

Yao Lu, Jiaxing QI, Zhongzhi Luan, Shaohan Huang, Carol Fung, Hailong Yang, and Depei Qian. FDLoRA: personalized federated learning of large language model via dual LoRA tuning, 2026. URL https://arxiv.org/abs/2406.07925.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT '11, pp. 142–150, USA, 2011. Association for Computational Linguistics. ISBN 9781932432879.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR, 20–22 Apr 2017. URL https://proceedings.mlr.press/v54/mcmahan17a.html.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 10 2019.

Raghav Singhal, Kaustubh Ponkshe, and Praneeth Vepakomma. FedEx-LoRA: Exact aggregation for federated and efficient fine-tuning of large language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1316–1336, Vienna, Austria, jul 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.67. URL https://aclanthology.org/2025.acl-long.67/.

Tao Sun, Dongsheng Li, and Bao Wang. Decentralized federated averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4289–4301, 2023. doi: 10.1109/TPAMI.2022.3196503.

Youbang Sun, Zitao Li, Yaliang Li, and Bolin Ding. Improving lora in privacy-preserving federated learning, 2024. URL https://arxiv.org/abs/2403.12313.

Alysa Ziying Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12):9587–9603, 2023. doi: 10.1109/TNNLS.2022.3160699.

Nazarii Tupitsa, Samuel Horváth, Martin Takáč, and Eduard Gorbunov. Selective collaboration for robust federated learning. In Rebekka Burkholz, Shiwei Liu, Saiprasad Ravishankar, William Redman, Wei Huang, Weijie Su, and Zhihui Zhu (eds.), *Conference on Parsimony and Learning*, volume 328 of *Proceedings of Machine Learning Research*, pp. 1161–1194. PMLR, 23–26 Mar 2026. URL https://proceedings.mlr.press/v328/tupitsa26a.html.

Ziyao Wang, Zheyu Shen, Yexiao He, Guoheng Sun, Hongyi Wang, Lingjuan Lyu, and Ang Li. Flora: Federated fine-tuning large language models with heterogeneous low-rank adaptations. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 22513–22533. Curran Associates, Inc., 2024. doi: 10.52202/079017-0708. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/28312c9491d60ed0c77f7fff4ad86dd1-Paper-Conference.pdf.

Jun Wu, Wenxuan Bao, Elizabeth Ainsworth, and Jingrui He. Personalized federated learning with parameter propagation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '23, pp. 2594–2605, New York, NY, USA, 2023a. Association for Computing Machinery. ISBN 9798400701030. doi: 10.1145/3580305.3599464. URL https://doi.org/10.1145/3580305.3599464.

Xinghao Wu, Xuefeng Liu, Jianwei Niu, Guogang Zhu, and Shaojie Tang. Bold but cautious: Unlocking the potential of personalized federated learning through cautiously aggressive collaboration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 19375–19384, October 2023b.

Xinghao Wu, Xuefeng Liu, Jianwei Niu, Haolin Wang, Shaojie Tang, and Guogang Zhu. FedloRA: When personalized federated learning meets low-rank adaptation, 2024. URL https://openreview.net/forum?id=bZh06ptG9r.

Jiuyun Xu, Liang Zhou, Yingzhi Zhao, Xiaowen Li, Kongshang Zhu, Xiangrui Xu, Qiang Duan, and RuRu Zhang. A two-stage federated learning method for personalization via selective collaboration. *Computer Communications*, 232:108053, 2025. ISSN 0140-3664. doi: https://doi.org/10.1016/j.comcom.2025.108053. URL https://www.sciencedirect.com/science/article/pii/S0140366425000106.

Yiyuan Yang, Guodong Long, Qinghua Lu, Liming Zhu, Jing Jiang, and Chengqi Zhang. Federated low-rank adaptation for foundation models: A survey. In James Kwok (ed.), *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pp. 10779–10787. International Joint Conferences on Artificial Intelligence Organization, 8 2025. doi: 10.24963/ijcai.2025/1196. URL https://doi.org/10.24963/ijcai.2025/1196. Survey Track.

Mang Ye, Xiuwen Fang, Bo Du, Pong C. Yuen, and Dacheng Tao. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Comput. Surv.*, 56(3), October 2023. ISSN 0360-0300. doi: 10.1145/3625558. URL https://doi.org/10.1145/3625558.

Liping Yi, Heng Zhang, Han Yu, Gang Wang, Xiaoguang Liu, and Qinghua Hu. pFedLoRA: model-heterogeneous personalized federated learning with homogeneous low-rank adapter sharing on mobile edge devices. *IEEE Transactions on Mobile Computing*, 25(9):13898–13911, 2026. doi: 10.1109/TMC.2026.3674996.

Xinzhi Yi, Chunqiang Hu, Bin Cai, Hongyu Huang, Yuwen Chen, and Kui Wang. FedALoRA: adaptive local LoRA aggregation for personalized federated learning in LLM. *IEEE Internet of Things Journal*, 12(24):51854–51865, 2025. doi: 10.1109/JIOT.2025.3582427.

Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging federated learning by local adaptation, 2022. URL https://arxiv.org/abs/2002.04758.

Michael Zhang, Karan Sapra, Sanja Fidler, Serena Yeung, and Jose M. Alvarez. Personalized federated learning with first order model optimization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=ehJqJQk9cw.

Wen Zhang, Lingfei Deng, Lei Zhang, and Dongrui Wu. A survey on negative transfer. *IEEE/CAA Journal of Automatica Sinica*, 10(2):305–329, 2023. ISSN 2329-9266. doi: 10.1109/JAS.2022.106004. URL https://www.ieee-jas.net/en/article/doi/10.1109/JAS.2022.106004.

Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, pp. 649–657, Cambridge, MA, USA, 2015. MIT Press.

Hao Zheng, Zhigang Hu, Liu Yang, Meiguang Zheng, Aikun Xu, and Boyu Wang. Confree: conflict-free client update aggregation for personalized federated learning. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'25/IAAI'25/EAAI'25, pp. 22875–22883. AAAI Press, 2025. ISBN 978-1-57735-897-8. doi: 10.1609/aaai.v39i21.34449. URL https://doi.org/10.1609/aaai.v39i21.34449.

A FULL BC-FEDLORA ALGORITHM

Algorithm 1 BC-FedLoRA

Input: Global LoRA parameters $\boldsymbol{\theta}^G$, clients set $\mathcal{N}$, maximal communication rounds T, and budget of participants K
Output: Final global LoRA parameters $\boldsymbol{\theta}^G_{T+1}$ and client-deployed parameters $\{\boldsymbol{\theta}_{i,T}\}_{i=1}^N$

Initialize participation counts $h_{i,1}$, personal states $\boldsymbol{\theta}^L_{i,0}$, deployed states $\boldsymbol{\theta}_{i,0}$, posterior statistics μ_0 and τ_0^2
for $t = 1, \ldots, T$ **do**
  Select the client subset $\mathcal{S}_t$ by equation 4
  for all $i \in \mathcal{S}_t$ **in parallel do**
    Determine the feasible LoRA rank cap $\boldsymbol{r}_i$ from client i's resource budget
    Train the federated branch with the current global model by equation 6
    Train the local proximal branch by equation 7
    Evaluate two branches and compute the benefit observation equation 8
    Conducting the Bayesian benefit update by equation 13 – equation 16
    Making the collaboration decision and deploying the model by equation 17
    Upload $(\boldsymbol{\Delta}^F_{i,t}, \boldsymbol{z}_i, \phi_i)$
    Update participation history $h_{i,t+1} \leftarrow h_{i,t} + 1$
  end for
  for all $i \notin \mathcal{S}_t$ **do**
    Set $\boldsymbol{\theta}_{i,t} \leftarrow \boldsymbol{\theta}_{i,t-1}$ and $h_{i,t+1} \leftarrow h_{i,t}$, preserve the posterior for these clients
  end for
  for all LoRA coordinates q **do**
    Update LoRA parameters by equation 18 and equation 19
  end for
end for
return $\boldsymbol{\theta}^G_{T+1}$ and $\{\boldsymbol{\theta}_{i,T}\}_{i=1}^N$

B ADDITIONAL EXPERIMENTAL RESULTS

B.1 ADDITIONAL MAIN RESULTS

BC-FedLoRA provides the highest mean accuracy in all Tiny-ImageNet-200 experiments, which can be observed in Table 5. The results in Figure 4 also shows the excellent convergence performance of our method. The accuracy of our method is still better than that of the strongest baseline, approximately being 1.1% and 0.7% percent under strong heterogeneity and approximately 0.2% percent in less-skewed scenarios.

Table 5: Mean client accuracy (%) on Tiny-ImageNet-200.

Method	$(\alpha = 0.1, \bar{r} = 1.5)$	$(\alpha = 0.1, \bar{r} = 2.5)$	$(\alpha = 0.5, \bar{r} = 1.5)$	$(\alpha = 0.5, \bar{r} = 2.5)$
Local-only	66.09 ± 0.88	65.89 ± 0.65	44.44 ± 0.96	44.43 ± 0.84
FedAvg-LoRA	68.87 ± 1.81	68.52 ± 1.40	70.43 ± 0.90	$\underline{70.70 \pm 1.42}$
FedProx-LoRA	68.99 ± 1.48	68.90 ± 0.88	$\underline{70.49 \pm 1.17}$	70.69 ± 1.08
FlexLoRA	66.55 ± 1.99	67.15 ± 1.27	69.60 ± 1.17	70.10 ± 1.45
FedEx-LoRA	66.83 ± 1.56	67.45 ± 1.33	70.10 ± 1.23	70.62 ± 1.36
Ditto-LoRA	$\underline{74.05 \pm 1.08}$	$\underline{75.24 \pm 0.80}$	66.03 ± 1.62	67.45 ± 1.26
FedFomo-LoRA	65.95 ± 0.36	65.79 ± 0.44	44.45 ± 0.91	44.16 ± 0.99
BC-FedLoRA	$\mathbf{75.16 \pm 0.84}$	$\mathbf{75.96 \pm 1.34}$	$\mathbf{70.66 \pm 0.59}$	$\mathbf{70.95 \pm 1.40}$

Meanwhile, BC-FedLoRA also ranks first in three settings of IMDB and second in the remaining one, as shown in Table 6. The gap to the strongest baseline is at 0.2% in most of settings.

When we look at the convergence curves in Figure 5, there is an interesting phenomenon that some of the baselines did not show obvious trend of convergence. Since the IMDB dataset only involves

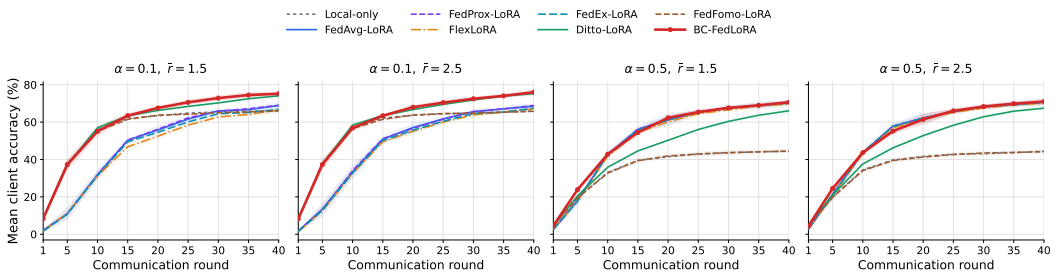

Figure 4: The convergence performance of all methods on Tiny-ImageNet-200.

Table 6: Mean client test accuracy (%) on IMDB.

Method	$(\alpha = 0.1, \bar{r} = 1.5)$	$(\alpha = 0.1, \bar{r} = 2.5)$	$(\alpha = 0.5, \bar{r} = 1.5)$	$(\alpha = 0.5, \bar{r} = 2.5)$
Local-only	93.42 ± 1.59	93.81 ± 1.24	88.31 ± 2.81	88.25 ± 3.24
FedAvg-LoRA	78.87 ± 7.89	78.90 ± 3.33	82.11 ± 1.86	81.25 ± 3.39
FedProx-LoRA	79.46 ± 7.20	78.12 ± 10.84	82.02 ± 2.01	80.88 ± 3.84
FlexLoRA	69.09 ± 8.78	67.57 ± 20.57	80.83 ± 2.92	79.61 ± 6.51
FedEx-LoRA	63.00 ± 13.43	72.34 ± 8.72	81.11 ± 3.01	79.80 ± 2.98
Ditto-LoRA	$\mathbf{94.83 \pm 1.54}$	$\underline{94.62 \pm 1.32}$	90.19 ± 2.40	$\underline{90.24 \pm 2.34}$
FedFomo-LoRA	94.49 ± 1.32	94.25 ± 1.09	$\underline{90.48 \pm 2.31}$	90.06 ± 2.68
BC-FedLoRA	$\underline{94.76 \pm 1.32}$	$\mathbf{94.82 \pm 1.38}$	$\mathbf{90.52 \pm 2.42}$	$\mathbf{90.44 \pm 2.12}$

two classes of label, the impact of heterogeneous clients and limited local samples on the direction of cross-round updates is quite pronounced. Fluctuations in label distributions and LoRA capacities among the clients can easily cause the global decision boundary to shift repeatedly. Since these baseline methods do not assess whether a specific sharing instance is beneficial, they may consistently aggregate updates that are mutually conflicting or even detrimental. Moreover, the wide confidence intervals (shade area in the figure) further indicate that their results are sensitive to data partitioning and training setups. In contrast, BC-FedLoRA estimates the gains of collaboration relative to local proximal adaptation and makes decisions based on Bayesian posteriors, thereby mitigating the persistent interference of unreliable updates on client models.

B.2 Detailed Collaboration-Decision Visualization

Figure 6 shows how often each collaboration state is occupied across checkpoints. The decision policy does not collapse to a single state. For CIFAR-100 and Tiny-ImageNet-200 with $\alpha = 0.5$, shared decisions account for roughly 85-88% of checkpoint states, consistent with the strong performance of federated baselines in these experiments. In contrast, IMDB retains 40-62% deferred states, and AG News also exhibits substantial deferral.

These patterns confirm that the learned policy produces heterogeneity- and resource-dependent decisions rather than a using a fixed scheme. These results are also consistent with the conditional text-task gains, because a large fraction of clients remain uncertain or personalized rather than contributing to a common model.

Overall, the experiments support BC-FedLoRA as a competitive personalized federated learning framework, it improves average and lower-tail accuracy in most tested conditions, while exposing when collaboration remains uncertain or harmful. To a certain extent, the negative-transfer issue can also be mitigated.

B.3 Ablation Discussions

In the following, we explain the observed mechanisms and failure modes. Particularly, HistInt replaces the Bayesian posterior with a historical sample interval. Latest uses only the latest observation and permits a decision after one observation. RandSel randomly selects clients. RandDep

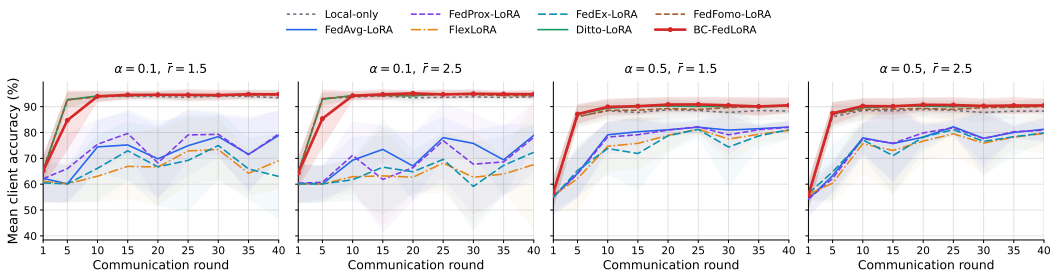

Figure 5: The convergence performance of all methods on IMDB.



Figure 6: Occupancy of shared, personalized, and deferred decisions across all recorded checkpoints.

randomly chooses the deployed branch. GATEUP uploads LoRA updates only after a shared decision, and LOCALREF jointly replaces the proximal personal branch and its comparator with a purely local reference. First of all, replacing history-balanced selection with uniform random selection reduces client's mean accuracy. Random deployment of adaptation branches is weaker than benefit-calibrated deployment. These comparisons indicate that the gain is not explained merely by how often clients share: both the clients from which evidence is collected and the mapping from evidence to deployment matter. Moreover, using the purely local reference degrades client's accuracy, and the results of GATEUP reveal that uploading a shared candidate only when the current deployment decision is shared decreases accuracy by $5.28/3.80$ on CIFAR-100 and $0.64/0.87$ on AG News, respectively. Therefore, deployment personalization should not be coupled with removing a selected client from global training. Withholding those candidates creates a feedback loop in which the shared model and the subsequent benefit evidence both deteriorate. Next, the two rule of evidence replacements HISTINT and LATEST also show the uniform advantage of Bayesian decision.