

---

# Federated Multi-armed Bandits with Personalization

---

Chengshuai Shi  
University of Virginia

Cong Shen  
University of Virginia

Jing Yang  
The Pennsylvania State University

## Abstract

A general framework of personalized federated multi-armed bandits (PF-MAB) is proposed, which is a new bandit paradigm analogous to the federated learning (FL) framework in supervised learning and enjoys the features of FL with personalization. Under the PF-MAB framework, a mixed bandit learning problem that flexibly balances generalization and personalization is studied. A lower bound analysis for the mixed model is presented. We then propose the Personalized Federated Upper Confidence Bound (PF-UCB) algorithm, where the exploration length is chosen carefully to achieve the desired balance of learning the local model and supplying global information for the mixed learning objective. Theoretical analysis proves that PF-UCB achieves an  $O(\log(T))$  regret regardless of the degree of personalization, and has a similar instance dependency as the lower bound. Experiments using both synthetic and real-world datasets corroborate the theoretical analysis and demonstrate the effectiveness of the proposed algorithm.

## 1 Introduction

Federated learning (FL) is an emerging distributed machine learning paradigm that has many attractive properties. In particular, FL is motivated by the growing trend that massive amounts of real-world data are exogenously generated at edge devices, which are non-independent and identically distributed (non-IID) and highly imbalanced (Bonawitz et al., 2019). FL focuses on many clients collaboratively training a machine learning model under the coordination of a cen-

tral server while keeping the local data private at each client (McMahan et al., 2017).

Earlier FL approaches focus on training a single global model that can perform well on the aggregated global dataset. However, the performance of the FL-trained global model on an individual client dataset degrades dramatically when significant heterogeneity among the local datasets exists, which raises the concern of using one global model for all individual clients in edge inference. To address this issue, FL with personalization (Smith et al., 2017) has been proposed. Instead of learning a single global model, each device aims at learning a mixture of the global model and its own local model (Hanzely and Richtárik, 2020; Deng et al., 2020), which provides an explicit trade-off between the two potentially competing learning goals.

While the main focus of the state of the art *FL with personalization* is on the supervised learning setting, we propose to extend its core principles to the multi-armed bandits (MAB) problem. This is motivated by a corpus of practical applications, including:

- **Cognitive radio.** Consider a cellular network where one base station (BS) serves many devices (e.g., smartphones) that are geographically spread out in the coverage area. Each device wants to use the individually best channel (in terms of its own communication quality) for data transmission, while the BS wants to learn a globally best channel averaged over the coverage area (e.g., to broadcast control information). Since BS is fixed at one location, the global channel quality cannot be measured by the BS itself – it has to come from measurements of the geographically distributed devices. However, if the devices make channel selection decisions only to learn the globally best channel, they may suffer from poor communication quality due to the local-global model mismatch. This calls for personalized federated MAB so that the global and local channel quality models are jointly considered.

- **Recommender system.** Local servers want to recommend the most popular items to their served customers to maximize the expected rewards. The item popularity can only be learned via interacting with customers, leading to a bandit problem (Li et al.,

2010). As different local servers have potentially heterogeneous customers, their local popularities are non-IID. In addition, each local server only collects data from a small group of customers, and the central server needs to average the locally learned popularity models to have a global model, without accessing the individual recommendation for privacy protection. However, the globally most popular item may not apply to the small group of customers of a particular local server, which again leads to personalization in a federated bandit setting, i.e., a joint consideration of global and local item popularities.

In both applications, the general FL characteristics need to be applied to an underlying bandit model, which greatly complicates the problem. The bandit setting is more difficult due to limited feedback (only observing one arm at a time) (Agarwal et al., 2020). In addition, FL has a deterministic pipeline, while the data collection for bandit is online and the server-clients coordination becomes dynamic. Moreover, incorporating personalization represents another significant challenge, since a client has to consider other clients using the already limited bandit feedback.

In this work, a novel framework of *personalized federated MAB (PF-MAB)* is developed, which represents a systematic attempt to bridge FL, MAB, and personalization. The PF-MAB framework generalizes the earlier works of federated bandits (Shi and Shen, 2021; Zhu et al., 2020; Dubey and Pentland, 2020) and can serve as an umbrella for a variety of bandit problems that share the FL principles and are in need of personalization. In particular, we claim the following contributions.

- A mixed global and local learning objective is studied in the PF-MAB framework, which allows for smoothly balancing *generalization* and *personalization* that depends on the specific application requirement.
- We provide a lower bound analysis of the general PF-MAB model, which reveals the fundamental requirement of balancing global and local explorations.
- Inspired by the lower bound analysis, we propose the *Personalized Federated Upper Confidence Bound (PF-UCB)* algorithm that carefully adjusts the lengths of local and global explorations based on the mixed learning objective. We also address the synchronization problem caused by client heterogeneity by leveraging the exploration-exploitation tradeoff.
- A rigorous regret upper bound analysis shows that PF-UCB achieves an  $O(\log(T))$  regret *regardless of the degree of personalization*, and has a similar instance dependency as shown in the lower bound.
- Additional algorithm enhancements guided by the

theoretical analysis are also discussed. We verify the effectiveness and efficiency of PF-UCB via numerical experiments on both synthetic and real-world datasets.

## 2 Problem Formulation

### 2.1 Single-player Stochastic MAB

In the standard stochastic MAB setting, a single player directly plays  $K$  arms, with rewards  $X_k$  of arm  $k \in [K]$  sampled independently from a  $\sigma$ -subgaussian distribution with mean  $\mu_k$ . At time  $t$ , the player chooses an arm  $\pi(t)$  and the goal is to maximize the expected cumulative reward in  $T$  rounds, i.e.,  $\mathbb{E} \left[ \sum_{t=1}^T X_{\pi(t)}(t) \right]$ , which is characterized by minimizing the regret:

$$R(T) = T\mu_* - \mathbb{E} \left[ \sum_{t=1}^T X_{\pi(t)}(t) \right], \quad (1)$$

where  $\mu_* := \mu_{k_*} = \max\{\mu_1, \dots, \mu_K\}$ . As shown by Lai and Robbins (1985), the regret is lower bounded by:

$$\liminf_{T \rightarrow \infty} \frac{R(T)}{\log(T)} \geq \sum_{k \neq k_*} \frac{\mu_* - \mu_k}{\text{kl}(X_k, X_{k_*})} \quad (2)$$

where  $\text{kl}(X_k, X_{k_*})$  denotes the KL-divergence between the two corresponding distributions.

### 2.2 PF-MAB

**Clients and local models.** In the PF-MAB framework, there are  $M$  clients interacting with the same set of  $K$  arms (referred as “local arms”). The clients are labeled from 1 to  $M$  to facilitate the discussion (labelling is not used in the algorithm). For client  $m$ , arm  $k$  generates local rewards  $X_{k,m}(t)$  independently from a  $\sigma$ -subgaussian distribution with mean  $\mu_{k,m}$ . Without loss of generality, we assume  $\sigma = 1$ . For different clients, their local models are non-IID, i.e., in general  $\mu_{k,n} \neq \mu_{k,m}$  when  $n \neq m$ . A client can only interact with her own local MAB model by choosing arm  $\pi_m(t)$  and receiving reward  $X_{\pi_m(t),m}(t)$  at time  $t$ . Also, there is no direct communication between clients.

**The global model.** A global stochastic MAB model with the same set of  $K$  arms (referred as “global arms”) coexists with the local models, where the global reward  $X_k(t)$  for the global arm  $k$  is the average of local rewards, i.e.,  $X_k(t) = \frac{1}{M} \sum_{m=1}^M X_{k,m}(t)$ . The global reward can be thought of as the *virtual* averaged reward had all  $M$  clients pulled the same arm  $k$  at time  $t$ . Correspondingly, the mean reward of global arm  $k$  is  $\mu_k = \frac{1}{M} \sum_{m=1}^M \mu_{k,m}$ . We note that although the global model is the average of local models, the global rewards are not directly observable by any client.

**Communication.** In decentralized multi-player multi-armed bandits (MP-MAB), clients are prohibited to have *explicit* communication with each other

(Liu and Zhao, 2010; Boursier and Perchet, 2019). We modify this constraint to enable client-server periodic communication that is similar to FL. Specifically, the clients can send “local model updates” to a central server, which then aggregates and broadcasts the updated “global model” to the clients. (We will specify these components later.) Note that just as in FL, communication is one of the major bottlenecks and the algorithm has to be conscious about its usage. This constraint is incorporated by imposing a loss  $C$  each time a communication round happens, which will be accounted for in the regret. We also make the assumption that clients and server are fully synchronized (McMahan et al., 2017).

## 2.3 Personalization vs Generalization

### 2.3.1 Two Extreme Cases

With the coexistence of local and global models, two extreme scenarios exist for the bandit learning: local-only and global-only. In the first case, clients only care about their own local performance, which is characterized by the local cumulative reward  $r_l(T)$  as

$$r_l(T) := \mathbb{E} \left[ \sum_{t=1}^T \sum_{m=1}^M X_{\pi_m(t),m}(t) \right].$$

$r_l(T)$  is equivalent to the sum rewards of  $M$  clients who play  $M$  decoupled and non-interacting MAB games. Obviously, the optimal choice for client  $m$  is arm  $k_{*,m}$  with  $\mu_{*,m} := \mu_{k_{*,m},m} = \max_{k \in [K]} \mu_{k,m}$ . However, only pursuing the locally optimal arm severely limits the ability of generalization across clients, especially when the degree of heterogeneity is significant.

For the other extreme, clients only focus on learning the global model, which means to maximize the global cumulative reward:

$$r_g(T) := \mathbb{E} \left[ \sum_{t=1}^T \sum_{m=1}^M X_{\pi_m(t)}(t) \right].$$

In this case, although the client’s action and observation are both on her local arms, the reward is defined with respect to the global arm (Shi and Shen, 2021). Ideally, the optimal choice to maximize  $r_g(T)$  is to let all the clients play the optimal global arm  $k_*$  with  $\mu_* := \mu_{k_*} = \max_{k \in [K]} \mu_k$ . We note that this problem is recently proposed and studied in Zhu et al. (2020); Shi and Shen (2021), which calls for efficient coordination among clients since no client can solve the global model individually. However, any efficient solution for this extreme case may lead to poor individual performance due to the non-IID local models.

### 2.3.2 Mixed Learning Objective

To balance the need of both personalization and generalization, we hereby introduce a new learning objective

which mixes  $r_g(T)$  and  $r_l(T)$  by a parameter  $\alpha \in [0, 1]$ . This learning objective is referred to as the *mixed cumulative reward*, which is defined as:

$$r(T) := \alpha r_l(T) + (1 - \alpha) r_g(T). \quad (3)$$

The parameter  $\alpha$  provides a flexible choice of personalization: with  $\alpha = 1$ ,  $r(T)$  becomes the sum rewards of  $M$  individual single-player MAB games (full personalization); with  $\alpha = 0$ ,  $r(T)$  only considers the global model (no personalization); with  $0 < \alpha < 1$ , both the global and local models are simultaneously taken into consideration by  $r(T)$ .

### 2.3.3 Mixed Model

An equivalent view of the mixed cumulative reward  $r(T)$  in Eqn. (3) is provided here, which facilitates our subsequent discussion. By unfolding  $r_l(T)$  and  $r_g(T)$ ,  $r(T)$  can be rewritten as

$$r(T) = \mathbb{E} \left[ \sum_{t=1}^T \sum_{m=1}^M X'_{\pi_m(t),m}(t) \right],$$

where  $X'_{\pi_m(t),m}(t)$  is a hypothetical reward that combines the local and global rewards, defined as:

$$X'_{\pi_m(t),m}(t) := \alpha X_{\pi_m(t),m}(t) + (1 - \alpha) X_{\pi_m(t)}(t). \quad (4)$$

Thus, maximizing the mixed cumulative reward can be equivalently viewed as playing a new MAB game with  $X'_{k,m}(t)$  as rewards for the clients. However, since clients cannot directly observe the global reward,  $X'_{k,m}(t)$  is only *partially observable* at each individual client. We refer to this hypothetical game as the *mixed model*. A similar reward definition using the weighted sum of clients’ rewards has been adopted in Brânzei and Peres (2019), albeit from a game theory perspective.

In client  $m$ ’s mixed model, the mean reward  $\mu'_{k,m} := \mathbb{E} [X'_{k,m}(t)]$  for arm  $k$  can be calculated as:

$$\mu'_{k,m} = \underbrace{\left( \alpha + \frac{1 - \alpha}{M} \right) \mu_{k,m}}_{\text{local info}} + \underbrace{\frac{1 - \alpha}{M} \sum_{n \neq m} \mu_{k,n}}_{\text{global info}}. \quad (5)$$

Since the global information in  $\mu'_{k,m}$  is determined by other clients and cannot be accessed directly at client  $m$ , communication between clients and the server is of critical importance.

With the mixed models, the notion of regret in Eqn. (1) can be generalized to  $r(T)$  as

$$R(T) = T \sum_{m=1}^M \mu'_{*,m} - \mathbb{E} \left[ \sum_{t=1}^T \sum_{m=1}^M X'_{\pi_m(t),m}(t) \right] + CMT_c, \quad (6)$$

where  $\mu'_{*,m}$  is the mean reward from the optimal arm  $k'_{*,m}$  of client  $m$ ’s mixed model, i.e.,  $\mu'_{*,m} := \mu'_{k'_{*,m},m} =$

$\max_{k \in [K]} \mu'_{k,m}$ . The first term in  $R(T)$  is the highest expected mixed cumulative reward that clients can get by always playing their optimal arms, which is similar to the optimal expected cumulative reward term  $T\mu_*$  in Eqn. (1). The additional loss term  $CMT_c$  in Eqn. (6) represents the communication loss, where  $T_c$  is the total amount of communication slots. Without loss of generality, we assume that the optimal arm of each client on her mixed model is unique. We further note that the optimal arms of different clients are likely to be different, because in general non-IID local models lead to  $k'_{*,m} \neq k'_{*,n}$  when  $m \neq n$ . We further denote  $\Delta'_{k,m} = \mu'_{*,m} - \mu'_{k,m}$ .

### 3 Lower Bound Analysis

A regret lower bound of PF-MAB is characterized in the following theorem.

**Theorem 1.** *For any consistent<sup>1</sup> algorithm  $\Pi$ , the regret  $R(T)$  in Eqn. (6) can be lower bounded as*

$$\liminf_{T \rightarrow \infty} \frac{R_\Pi(T)}{\log(T)} \geq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \max \left\{ \frac{\Delta'_{k,m}}{\text{kl}(Y_{k,m}, Y_{k'_{*,m},m})}, \frac{\Delta'_{k,m}}{\min_{n:n \neq m, k'_{*,n} \neq k} \text{kl}(Z_{k,n}^m, Z_{k'_{*,n},n}^m)} \right\}, \quad (7)$$

where  $Y_{k,m} = (\alpha + \frac{1-\alpha}{M}) X_{k,m} + \mu'_{k,m} - (\alpha + \frac{1-\alpha}{M}) \mu_{k,m}$  and  $Z_{k,n}^m = \frac{1-\alpha}{M} X_{k,m} + \mu'_{k,n} - \frac{1-\alpha}{M} \mu_{k,m}$ .

The proof of Theorem 1 can be found in the supplementary material. The communication cost is ignored in the analysis (i.e.,  $C = 0$ ), but naturally, this lower bound still holds for  $C > 0$ . The lower bound in Eqn. (7) sums over the maximum of two terms for all clients and suboptimal arms. First, random variable  $Y_{k,m}$  with mean  $\mu'_{k,m}$  represents an idealized degenerated game of client  $m$ 's mixed model where information from other clients, i.e.,  $\{\mu_{k,l}\}_{l \neq m}$ , is perfectly known. With  $Y_{k,m}$ , a lower bound for the regret of client  $m$  learning arm  $k$  for her mixed model can be obtained. Second, random variable  $Z_{k,n}^m$  with mean  $\mu'_{k,n}$  represents another idealized degenerated game of client  $n$ 's mixed model, where we assume full information of arm  $k$  from all other clients except client  $m$ , i.e.,  $\{\mu_{k,l}\}_{l \neq m}$ . With  $Z_{k,n}^m$ , the regret of client  $m$  providing information of arm  $k$  to client  $n$  is characterized. Then, building on this characterization, the regret of client  $m$  providing information of arm  $k$  to all other clients can be lower bounded by taking the worst case among them, i.e., the minimization term. This worst-case argument corresponds to the client who requires the most global information of arm  $k$ . To summarize,

<sup>1</sup>The consistent algorithm is defined the same way as in Lai and Robbins (1985) but with the regret of Eqn. (6).

the first and second terms in the maximization characterize the necessary loss for learning local (for the client herself) and global (for all other clients) information of client  $m$ 's arm  $k$ , respectively. We also note that in the case of  $\alpha = 1$ , i.e., local-only, Eqn. (7) recovers the lower bound in Eqn. (2), summed over  $M$  local models.

More light can be shed on the lower bound by limiting the attention to Gaussian distributed rewards.

**Corollary 1.** *For any consistent algorithm  $\Pi$ , if the rewards follow Gaussian distributions with unit variance, the regret is lower bounded as*

$$\liminf_{T \rightarrow \infty} \frac{R_\Pi(T)}{\log(T)} \geq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \max \left\{ \frac{2\beta^2}{\Delta'_{k,m}}, \frac{2\gamma^2 \Delta'_{k,m}}{(\Delta'_k)^2} \right\},$$

where  $\beta = \alpha + \frac{1-\alpha}{M}$ ,  $\gamma = \frac{1-\alpha}{M}$  and  $\Delta'_k = \min_{n:k'_{*,n} \neq k} \Delta'_{k,n}$ .

Corollary 1 shows that the second term in the maximum is determined by  $\Delta'_k$ , which corroborates that the loss of learning global information for arm  $k$  is determined by the hardest mixed model.

We note that, as will be evident in the PF-UCB algorithm, the lower bound analysis reveals important guidelines for balancing global and local explorations. Nevertheless, neither Theorem 1 nor Corollary 1 establishes a *universally tight* lower bound (for all  $\alpha$ ). Characterizing the precise lower bound dependency on  $\alpha$  is an interesting open problem, and we have the following conjecture.

**Conjecture 1.** *For any consistent algorithm  $\Pi$ , as  $T \rightarrow \infty$ ,  $\forall m \in [M]$  and  $\forall k : k \neq k'_{*,m}$ , it holds that*

$$\frac{\beta^2}{T_{k,m}} + \sum_{n:n \neq m, k'_{*,n} \neq k} \frac{\gamma^2}{T_{k,n}} \leq \frac{\eta^2 \text{kl}(X'_{k,m}, X'_{k'_{*,m},m})}{\log(T)},$$

where  $T_{k,m}$  is the expected number of pulls on arm  $k$  by client  $m$  in the  $T$  time slots,  $\beta = \alpha + \frac{1-\alpha}{M}$ ,  $\gamma = \frac{1-\alpha}{M}$  and  $\eta = (\beta^2 + (M-1)\gamma^2)^{\frac{1}{2}}$ .

Conjecture 1 also recovers Eqn. (2) with  $\alpha = 1$ . Furthermore, with  $\alpha = 0$  (global-only), it implies that  $\liminf_{T \rightarrow \infty} \frac{R(T)}{\log(T)} \geq \sum_{k \neq k_*} \frac{M\Delta_k}{\text{kl}(X_k, X_{k_*})}$ , where  $\Delta_k = \mu_* - \mu_k$ . This result is reasonable as it is equivalent to the lower bound of a centralized client who directly maximizes the cumulative global reward.

## 4 PF-UCB Algorithm

### 4.1 Challenges

Solving the PF-MAB model faces several new challenges. The first challenge is that in order to maximize

the mixed reward, both local and global information are essential. On one hand, the overall decision can be compromised (depending on the choice of  $\alpha$ ) as long as one type of information is insufficiently learned. On the other hand, providing global information for other clients may degrade the individual performance since the additional exploration does not directly benefit the client. The key challenge is how to gain *sufficient but not excessive* local and global information simultaneously based on the required degree of personalization.

A second challenge is that the game difficulties vary across clients. It is highly likely that different clients would need different amounts of global information. In other words, some clients may find their optimal arms much slower than the others, which is similar to the client heterogeneity problem in FL (Li et al., 2020). How to handle the resulting synchronization problem caused by client heterogeneity in PF-MAB becomes an important issue.

Lastly, although communication is fundamental to providing global information, it incurs additional losses in regret. This benefit-cost balance needs to be addressed in the algorithm design.

## 4.2 Algorithm Design

The Personalized Federated Upper Confidence Bound (PF-UCB) algorithm operates in phases (analogous to communication rounds in FL), and each phase consists of three sub-phases: global exploration, local exploration, and exploitation. The set of arms for global (resp. local) exploration are referred to as the set of global (resp. local) active arms. Specifically, at phase  $p$ ,  $A_m(p)$  (with cardinality  $K_m(p)$ ) and  $A(p) = \cup_{m \in [M]} A_m(p)$  (with cardinality  $K(p)$ ) denote the set of local and global active arms respectively, which are both initialized as  $[K]$ . PF-UCB for clients and the central server are presented in Algorithms 1 and 2, respectively.

In phase  $p$ , global exploration is first performed in order to collect statistics to update the global information. Client  $m$  explores each arm  $k \in A(p)$ , i.e., global active arms, for  $n_{k,m}^g(p) = \lceil (1 - \alpha)f(p) \rceil$  times, and thus the entire global exploration sub-phase lasts for  $K(p)\lceil (1 - \alpha)f(p) \rceil$  time slots. Note that  $f(p)$  is a flexible exploration length determined by the phase index  $p$ , and its impact on the regret is analyzed later. Since all clients share the same global active arm set  $A(p)$ , the global exploration length is also the same for them.

After the global exploration, the clients perform local exploration to update the local information. Each arm  $k \in A_m(p)$  is played by client  $m$  for  $n_{k,m}^l(p) = \lceil M\alpha f(p) \rceil$  times, which means the local exploration lasts for  $K_m(p)\lceil M\alpha f(p) \rceil$  time slots at client  $m$ . It is

---

### Algorithm 1 PF-UCB: client $m$

---

**Input:**  $T, M, K, \alpha$

- 1: Initialize  $p \leftarrow 1; A_m(1), A(1) \leftarrow [K]; \forall k \in [K], s_{k,m} \leftarrow 0, T_{k,m} \leftarrow 0; g, h \leftarrow 1; O_m \leftarrow 0$
- 2: **while**  $A(p) \neq \emptyset$  **do**
  - ▷ **Global exploration**
  - 3: **for**  $g \leq K(p) \lceil (1 - \alpha)f(p) \rceil$  **do**
  - 4:  $\pi \leftarrow (g \bmod K(p))$ -th arm in  $A(p)$
  - 5: Pull arm  $\pi$  and receive reward  $r_\pi$
  - 6:  $s_{\pi,m} \leftarrow s_{\pi,m} + r_\pi; T_{\pi,m} \leftarrow T_{\pi,m} + 1; g \leftarrow g + 1$
  - 7: **end for**
  - ▷ **Local exploration**
  - 8: **for**  $h \leq K_m(p) \lceil M\alpha f(p) \rceil$  **do**
  - 9:  $\pi \leftarrow (h \bmod K_m(p))$ -th arm in  $A_m(p)$
  - 10: Pull arm  $\pi$  and receive reward  $r_\pi$
  - 11:  $s_{\pi,m} \leftarrow s_{\pi,m} + r_\pi; T_{\pi,m} \leftarrow T_{\pi,m} + 1; h \leftarrow h + 1$
  - 12: **end for**
  - 13: Update  $\bar{\mu}_{k,m}(p) \leftarrow s_{k,m}/T_{k,m}, \forall k \in A(p)$
  - 14: Send  $\bar{\mu}_{k,m}(p), \forall k \in A(p)$  to the server
  - ▷ **Exploitation**
  - 15: **if**  $O_m = 0$  **then**
  - 16:  $k_{*,m}^l(p) \leftarrow \arg \max_{k \in A_m(p)} \{\bar{\mu}'_{k,m}(p - 1)\}$
  - 17: **else**  $k_{*,m}^l(p) \leftarrow O_m$
  - 18: **end if**
  - 19: Pull arm  $k_{*,m}^l(p)$  until receiving  $\bar{\mu}_k(p), k \in A(p)$
  - 20:  $\forall k \in A_m(p), \bar{\mu}'_{k,m}(p) \leftarrow \alpha \bar{\mu}_{k,m}(p) + (1 - \alpha) \bar{\mu}_k(p)$
  - 21: Update  $E_m(p)$  as in Eqn. (8) ▷ *Arm elimination*
  - 22:  $A_m(p + 1) \leftarrow A_m(p) \setminus E_m(p)$
  - 23: **if**  $|A_m(p + 1)| = 1$  **then**
  - 24:  $O_m \leftarrow$  the only arm in  $A_m(p + 1); A_m(p + 1) \leftarrow \emptyset$
  - 25: **end if**
  - 26: Send  $A_m(p + 1)$  to the server
  - 27: Receive  $A(p + 1)$  from the server;  $p \leftarrow p + 1; g, h \leftarrow 1$
  - 28:  $K(p + 1) \leftarrow |A(p + 1)|, K_m(p + 1) \leftarrow |A_m(p + 1)|$
  - 29: **end while**
  - 30: Stay on arm  $O_m$  until  $T$  ▷ *Exploitation*

---

important to note that since different clients may have local exploration sets of different sizes, i.e.,  $K_m(p)$  can be different across  $m$ , the local exploration length may also vary across clients.

Note that the lengths of global and local explorations are carefully designed. For each arm  $k \in A_m(p)$ , it is explored for  $\lceil (1 - \alpha)f(p) \rceil$  times during global exploration (recall that  $A_m(p) \subseteq A(p) = \cup_{m \in [M]} A_m(p)$ ) and  $\lceil M\alpha f(p) \rceil$  times during local exploration, leading to a total of  $n_{k,m}(p) = \lceil (1 - \alpha)f(p) \rceil + \lceil M\alpha f(p) \rceil$  pulls by client  $m$ . At the same time, client  $m$  is also assured that arm  $k$  is pulled by every other client  $n$  for at least  $n_{k,n}^g(p) = \lceil (1 - \alpha)f(p) \rceil$  times since they share the same  $A(p)$ . Thus, the proportion between local and global information is  $\frac{(1-\alpha)+M\alpha}{(1-\alpha)}$ , which coincides with the desired allocation in Eqn. (5)<sup>2</sup>.

After completing both global and local explorations, client  $m$  first sends the updated local sample means

---

<sup>2</sup>A detailed discussion on matching  $n_{k,m}(p)/n_{k,n}^g(p)$  with the weight ratio of Eqn. (5) is provided in the supplementary material.

---

**Algorithm 2** PF-UCB: central server
 

---

**Input:**  $T, M, K$   
 1: Initialize  $p \leftarrow 1$ ;  $A(1) \leftarrow [K]$   
 2: **while**  $A(p) \neq \emptyset$  **do**  
 3:   Receive  $\bar{\mu}_{k,m}(p), \forall k \in A(p)$  from all clients  $m \in [M]$   
 4:   Update  $\bar{\mu}_k(p) \leftarrow \frac{1}{M} \sum_{m=1}^M \bar{\mu}_{k,m}(p), \forall k \in A(p)$   
 5:   Send  $\bar{\mu}_k(p), \forall k \in A(p)$  to all clients  
 6:   Receive  $A_m(p+1)$  from all clients  
 7:   Send  $A(p+1) \leftarrow \cup_{m \in [M]} A_m(p+1)$  to all clients  
 8:    $p \leftarrow p+1$   
 9: **end while**

---

of all global active arms  $k \in A(p)$ , denoted as  $\bar{\mu}_{k,m}(p)$  for arm  $k$  at phase  $p$ , as the ‘‘local model updates’’ to the server. Since the local exploration length may vary, the server may not receive the updates from all clients at the same time. Thus, it has to wait until the updated sample means from all the clients are received and then sends the aggregated ‘‘global model’’  $\bar{\mu}_k(p) = \frac{1}{M} \sum_{m=1}^M \bar{\mu}_{k,m}(p)$  back to the clients. While this waiting time is necessary to synchronize the clients, it also leads to an increased regret, i.e., all clients have to wait for the slowest client before the next iteration.

In PF-MAB, The celebrated exploration-exploitation tradeoff in MAB is embraced to keep the regret caused by this waiting time low. The idea is that clients who have already sent local updates can begin *exploitation* while the server still waits to collect information from other clients. Specifically, before  $\bar{\mu}_k(p)$  are sent back, client  $m$  keeps playing her empirically best arm  $\bar{k}'_{*,m} = \arg \max_{k \in A_m(p)} \{\bar{\mu}'_{k,m}(p-1)\}$ , where  $\bar{\mu}'_{k,m}(p-1)$  is the estimation of  $\mu'_{k,m}$  in the preceding phase. Regret analysis shows that this is essential in keeping clients update periodically synchronized while achieving a low regret.

After the global sample means  $\bar{\mu}_k(p)$  are broadcast to the clients, the estimation for  $\mu'_{k,m}$  is updated as  $\bar{\mu}'_{k,m}(p) = \alpha \bar{\mu}_{k,m}(p) + (1-\alpha) \bar{\mu}_k(p)$ . Then, a local arm elimination procedure is performed such that the arms that are sub-optimal with a high probability are eliminated. With the newly calculated  $\bar{\mu}'_{k,m}(p)$ , the elimination set  $E_m(p)$  can be constructed as:

$$\left\{ k : k \in A_m(p), \max_{l \in A_m(p)} \bar{\mu}'_{l,m}(p) - \bar{\mu}'_{k,m}(p) \geq 2B_p \right\}, \quad (8)$$

where  $B_p = \sqrt{4 \log(T) / (MF(p))}$  is the confidence bound and  $F(p) = \sum_{q=1}^p f(q)$ . Note that the simple and clean form of  $B_p$  comes from the carefully designed lengths of global and local explorations. The local active set  $A_m(p+1)$  for the next phase is updated as  $A_m(p+1) = A_m(p) \setminus E_m(p)$ . Finally, all the clients send  $A_m(p+1)$  to the server and subsequently receive the global active set  $A(p+1) = \cup_{m \in [M]} A_m(p+1)$  from the server. As long as an arm is in the local active set

of at least one client, it is contained in the global active set because more global information regarding this arm is still needed to help (at least) that client make decisions.

When the local active set contains only one arm, i.e.,  $|A_m(q)| = 1$ , client  $m$  marks the only left arm in  $A_m(q)$  as the fixed arm  $O_m$  and sets  $A_m(q) = \emptyset$ . Then, she only sends an empty set to the server for the local active set update since her optimal arm is found. Also, with  $A_m(q) = \emptyset$ , client  $m$  does not perform local explorations any more. Nevertheless, global exploration is still necessary for client  $m$  as long as  $A(q)$  is not empty, because other clients still need information from her. In the exploitation phase, she also directly plays the fixed arm  $O_m$ . When all clients have found their optimal arms, i.e.,  $A(q) = \emptyset$ , they all fixate on their identified arms until the end of  $T$  without any further communication.

**Remarks.** It can be observed that the choice of local exploration length scales linearly with the number of clients, i.e.,  $n_{k,m}^l(p) = \lceil M \alpha f(p) \rceil \propto M$ , which may not be desirable when  $M$  is large. It is possible to simultaneously scale down the local and global exploration lengths as  $M$  increases, e.g.,  $n_{k,m}^l(p) = \lceil \alpha f(p) \rceil$  and  $n_{k,m}^g(p) = \lceil (1-\alpha) f(p) / M \rceil$ , to further trade off exploration and communication, but this does not fundamentally change the regret behavior that is to be discussed. A final note is that only sample means and active sets are communicated in the entire procedure – no raw samples and number of pulls are shared. This is similar to sharing model updates instead of raw data samples in FL, which helps preserve privacy.

## 5 Regret Analysis

The theoretical analysis for PF-UCB is presented in this section. In particular, Theorem 2 characterizes a regret upper bound of PF-UCB.

**Theorem 2.**  $\forall m \in [M]$  and  $\forall k \neq k'_{*,m}$ , suppose  $p'_{k,m}$  is the smallest integer that satisfies

$$MF(p'_{k,m}) \geq \frac{64 \log(T)}{(\Delta'_{k,m})^2}. \quad (9)$$

The regret of PF-UCB can be bounded as

$$\begin{aligned} R(T) &\leq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_{k,m}} [\alpha M f(p)] \\ &+ \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_k} [(1-\alpha) f(p)] \\ &+ \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_{k,m}} K [\alpha M f(p)] P'_{k,m}(p) \\ &+ 2CMp'_{\max} + 2(1+2C)M^2K, \end{aligned} \quad (10)$$

where  $p'_k = \max_{m \in [M]} \{p'_{k,m}\}$ ,  $p'_{\max} = \max_{k \in [K]} \{p'_k\}$  and  $P'_{k,m}(p) = \exp\{-\Delta'^2_{k,m} MF(p-1) / 4\}$ .

Detailed proof of Theorem 2 can be found in the supplementary material, which shows that the total regret can be decomposed into local and global exploration losses, exploitation loss, and communication loss. Note that the local exploration loss (the first term) is determined individually by each client's local model, i.e.,  $p'_{k,m}$ , while the global exploration loss (the second term) is determined globally, i.e.,  $p'_k$ . This coincides with Theorem 1 and Corollary 1. In addition, there is no global (resp. local) exploration loss in the local-only (resp. global-only) scenario, i.e.,  $\alpha = 1$  (resp. 0). Furthermore, although the constant term  $2(1+2C)M^2K$  in Eqn. (10) has a dependence on  $M^2$ , one may trade off this term with other regret terms by adjusting the confidence bound, e.g., specifying  $B_p = \sqrt{4\log(MT)/(MF(p))}$ .

While Theorem 2 provides a general characterization with unspecified  $f(p)$ , the following corollary gives an explicit form of regret with  $f(p) = 2^p \log(T)$ .

**Corollary 2.** *With  $f(p) = 2^p \log(T)$ , it holds that*

$$R(T) = O\left(\sum_{m=1}^M \sum_{k \neq k'_{*,m}} \left[ \frac{\alpha}{\Delta'_{k,m}} + \frac{1-\alpha}{M} \frac{\Delta'_{k,m}}{(\Delta'_k)^2} \right] \log(T)\right),$$

where  $\Delta'_k = \min_{n:k'_{*,n} \neq k} \{\Delta'_{k,n}\}$ .

With this choice, PF-UCB achieves an  $O(\log(T))$  regret *regardless of*  $\alpha$ . It also has a similar instance dependency on  $\Delta'_{k,m}$  and  $\Delta'_k$  as shown in Corollary 1. Although the  $\alpha$ -dependency does not match Corollary 1, which is not necessarily a tight lower bound, Corollary 2 does match the sum of single-player lower bound when  $\alpha = 1$ . Interestingly, when  $\alpha = 0$ , the achievable upper bound in Corollary 2 approaches the conjectured lower bound in Conjecture 1. It is also worth noting that the communication and exploitation losses when  $f(p) = 2^p \log(T)$  are both of order  $O(1)$ , which demonstrates its efficiency. Regrets with other choices of  $f(p)$  can be found in the supplementary material.

We highlight the key components in the proof of Theorem 2 and Corollary 2 in the remainder of this section. A typical event

$$G = \{|\bar{\mu}'_{k,m}(p) - \mu'_{k,m}| \leq B_p, \forall p, \forall m \in [M], \forall k \in A_m(p)\}$$

is first established, and we can show that event  $G$  happens with high probability.

**Lemma 1.** *It holds that  $\mathbb{P}(G) := P_G \geq 1 - \frac{2MK}{T}$ .*

We then analyze the different loss components of the total regret in the following.

### 5.1 Exploration Loss

First, Lemma 2 bounds the number of pulls at clients on their sub-optimal arms.

**Lemma 2.** *Suppose event  $G$  happens. For client  $m$ , sub-optimal arm  $k \neq k'_{*,m}$  is guaranteed to be eliminated by phase  $p'_{k,m}$  as defined in Theorem 2.*

Then, the local and global exploration losses, denoted as  $R_l^{expr}(T)$  and  $R_g^{expr}(T)$ , respectively, can be bounded by the following lemma.

**Lemma 3.** *Suppose event  $G$  happens. With  $p'_{k,m}$  and  $p'_k$  defined in Theorem 2,  $R_l^{expr}(T)$  and  $R_g^{expr}(T)$  can be bounded, respectively, as*

$$R_l^{expr}(T) \leq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_{k,m}} [\alpha M f(p)],$$

$$R_g^{expr}(T) \leq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_k} [(1-\alpha)f(p)].$$

Note that  $R_g^{expr}(T)$  for arm  $k$  is determined by  $p'_k$ , which is from the hardest local model for arm  $k$ . It also matches Theorem 1 and Corollary 1.

### 5.2 Exploitation Loss

The exploitation loss  $R^{expt}(T)$  is caused by the exploitations when a client has to wait for other clients. Noting that this loss stops once the optimal arm is declared,  $R^{expt}(T)$  can be bounded as follows.

**Lemma 4.** *Suppose event  $G$  happens. With  $p'_{k,m}$  and  $P'_{k,m}(p)$  defined in Theorem 2,  $R^{expt}(T)$  can be bounded as*

$$R^{expt}(T) \leq \sum_{m=1}^M \sum_{k \neq k'_{*,m}} \Delta'_{k,m} \sum_{p=1}^{p'_{k,m}} K [M\alpha f(p)] P'_{k,m}(p).$$

### 5.3 Communication Loss

Since communication stops once all the optimal arms are declared, the communication loss is bounded as:

**Lemma 5.** *Suppose event  $G$  happens. With  $p'_{\max}$  defined in Theorem 2, the communication loss  $R^{comm}(T)$  can be bounded as*

$$R^{comm}(T) \leq 2CMp'_{\max}.$$

With Lemmas 1 to 5, Theorem 2 can be proved.

## 6 Algorithm Enhancement

While the exploration length in Section 4 can be viewed as evenly splitting the workload among clients (especially for the global exploration), it ignores the fact that the same action results in different losses at




Figure 1: Synthetic Regret




Figure 2: MovieLens Regret




Figure 3: MovieLens Reward




Figure 4: Enhancement

different clients. We propose an enhancement to adaptively adjust the exploration lengths for client  $m$ , as follows:

$$n_{k,m}^l(p) \propto \frac{\alpha M f(p)}{(\Delta'_{k,m})^{1/2}}, \forall k \in A_m(p), k \neq k'_{*,m};$$

$$n_{k,m}^g(p) \propto \frac{(1-\alpha)f(p)}{(\Delta'_{k,m})^{1/2}}, \forall k \in A(p), k \neq k'_{*,m}.$$

More details on designing this enhancement can be found in the supplementary material. Note that the exploration length for arm  $k$  is now proportional to  $1/(\Delta'_{k,m})^{1/2}$ , which coincides with the intuition that the workload should decrease for those clients who suffer large losses, i.e., with large  $\Delta'_{k,m}$ 's. However, this is difficult to implement without the knowledge of  $\Delta'_{k,m}$ . One way to resolve this is to assume all of the sub-optimal gaps are the same, which results in the chosen length in Section 4. In this enhancement, however, we propose to replace  $\Delta'_{k,m}$  by an estimation  $\bar{\Delta}'_{k,m}(p)$  in phase  $p$ , which can be specified as

$$\bar{\Delta}'_{k,m}(p) = \max_{l \in [K]} \bar{\mu}'_{l,m}(p-1) - \bar{\mu}'_{k,m}(p-1) + 2B_{p-1}.$$

Rigorously analyzing the regret of this enhancement turns out to be difficult, and we evaluate it only through experiments.

## 7 Experiments

Experiment results using both synthetic and real-world datasets are reported in this section to evaluate PF-UCB and the proposed enhancement. The communication loss is set as  $C = 1$  and  $f(p)$  is set to be  $2^p \log(T)$ . Details of the experiments (including the implementation codes) and additional results can be found in the supplementary material.

First, PF-UCB is evaluated with various choices of  $\alpha$  under a synthetic bandit game with 4 clients and 9 arms. The game is carefully designed such that all clients have different local optimal arms and the global optimal arm is also sub-optimal locally. Fig. 1 shows that PF-UCB successfully converges to the optimal choices across different values of  $\alpha$ , which proves its effectiveness in handling different combinations of personalization and generalization. The varying overall regrets and convergence speeds are the result of different game difficulties associated with different  $\alpha$ .

We then return to one of the motivating examples – the recommender system – and utilize the real-world MovieLens dataset (Cantador et al., 2011) for an empirical study of PF-MAB. The 2113 clients and 10197 movies in the dataset are randomly divided into 10 and 40 groups, respectively, and the averaged movie ratings from each group of clients are used to construct their local rewards, which vary across the groups of clients and naturally lead to non-IID local models. This game is larger and harder than the previous synthetic game. Especially, some groups have suboptimality gaps on their mixed models at around  $10^{-4}$ . As shown in Fig. 2, sub-linear regrets are achieved by PF-UCB with different values of  $\alpha$ . Note that in some cases (e.g.,  $\alpha = 0.1$ ), the algorithm does not completely converge within the given horizon; however, the regret curve only increases slowly at the end, which suggests that most of the suboptimal arms are eliminated.<sup>3</sup>

We also evaluate the rewards instead of regrets in the same MovieLens experiments. Fig. 3 reports the averaged per-step reward that PF-UCB achieves with varying  $\alpha$ . The optimal global and local rewards (labeled as “best global” and “best local”) represent the theoretically highest global and local mean rewards, respectively. The mixed, global, and local rewards (labeled as “mixed”, “global”, and “local”) generated by the actions of clients with PF-UCB are plotted. Fig. 3 shows that the mixed and global rewards almost meet the optimal global rewards with  $\alpha = 0$  (global-only), while the local rewards are highly sub-optimal. With an increase of  $\alpha$ , the mixed and local rewards trend up, indicating the focus is gradually shifted towards the local rewards, and simultaneously the global rewards trend down. At  $\alpha = 1$  (local-only), the mixed and local rewards almost achieve the optimal local rewards, while the global rewards are poor. This gradual shifting shows that introducing  $\alpha$  provides a smooth tradeoff between local and global rewards.

Lastly, the algorithm enhancement in Section 6 is evaluated. With a 4-client 9-arm game, the performance of the original and enhanced PF-UCB is compared in Fig. 4 with  $\alpha = 0.5$ . It can be observed that both algo-

<sup>3</sup>We note that in practice, the dataset is likely to be structured more carefully, e.g., grouping movies by categories instead of randomly, which would in general lead to easier games and faster convergence.



gorithms converge but the enhanced design has a lower regret, demonstrating its effectiveness.

## 8 Discussions

The proposed PF-MAB framework, the PF-UCB algorithm, and their companion theoretical analysis represent our initial attempt to bridge FL, MAB, and personalization. It also leaves several interesting directions for future research.

First, PF-UCB falls in the category of “cross-silo” federated learning (Kairouz et al., 2019), where clients are fixed, reliable and always available. Another important category of FL is the “cross-device” model, where only a fraction of clients are available at one time and some of them may fail or drop out. It is interesting to study the PF-MAB framework in such cross-device setting. Some simple modifications can be made to PF-UCB for this setting, e.g., sampling from the randomly available clients as participants. It is however challenging to have a rigorous theoretical analysis with time-varying participating clients.

Second, as stated in Section 3, it would be of great value to have a *tight* lower bound analysis under the PF-MAB model, and (dis)proving Conjecture 1 may serve as a starting point. We believe the main difficulty comes from that, instead of the traditional exploration-exploitation tradeoff, learning in PF-MAB actually faces a much more complicated tradeoff among global exploration, local exploration, and exploitation. Moreover, communication loss is ignored in Theorem 1 and Conjecture 1, and it would further complicate the tradeoff if communication loss is considered.

Last but not the least, as the algorithm enhancements proposed in Section 6 are validated only through experiment, it would be valuable to have rigorous analysis on whether (and how much) it outperforms the original PF-UCB.

## 9 Related Works

**Differences to FL.** FL has been an active research area over the past few years (McMahan et al., 2017; Bonawitz et al., 2019), with many attractive features as discussed in Section 1. See Kairouz et al. (2019); Li et al. (2020) for comprehensive surveys of FL. In particular, FL with personalization (Kulkarni et al., 2020) is an emerging topic, where learning a mixed local and global model (as in PF-MAB) is a representative approach (Hanzely and Richtárik, 2020; Deng et al., 2020; Mansour et al., 2020) among others (Wang et al., 2019; Smith et al., 2017; Jiang et al., 2019; Fal-

lah et al., 2020). Nevertheless, existing studies on FL are almost exclusively on supervised learning and there is very limited literature considering bandit (Shi and Shen, 2021; Zhu et al., 2020).

**Differences to Multi-player MAB.** The decentralized MP-MAB problem is related to PF-MAB but fundamentally different. The MP-MAB research considers either the “cooperative” setting (Landgren et al., 2016, 2018; Wang et al., 2020) or the “competitive” setting (Rosenski et al., 2016; Boursier and Perchet, 2019; Shi et al., 2020). Although user-dependent local models are studied in both settings (Shahrampour et al., 2017; Bistriz and Leshem, 2018; Boursier et al., 2020), our work is the first to study a flexible and mixed learning objective with partially observable rewards in MAB, to the best of our knowledge.

**Recent Advances.** A few recent works have touched upon the concept of federated bandits but none of them systemically addresses the key challenges introduced by personalization. Li et al. (2020); Dubey and Pentland (2020) assume IID local models and focus on privacy protection. Agarwal et al. (2020) studies regression-based contextual bandits as an example of the federated residual learning framework, which does not generalize to our formulation. The recent studies in Zhu et al. (2020); Shi and Shen (2021) are more related to this work, where federated MAB without personalization (i.e., global-only) is studied. A similar client-server communication protocol is adopted in Shi and Shen (2021) while a gossiping information-sharing strategy is applied in Zhu et al. (2020).

## 10 Conclusions

In this work, we have developed a general PF-MAB framework to bridge MAB, FL, and personalization. By focusing on learning a mixed global-local objective, this framework enables a flexible tradeoff between personalization and generalization. A lower bound analysis was provided for PF-MAB. The proposed PF-UCB algorithm caters to the need of personalization and addresses client heterogeneity by leveraging the exploration-exploitation tradeoff. Theoretical analysis showed that PF-UCB can achieve a regret of  $O(\log(T))$  regardless of the degree of personalization, and share a similar instance dependency as the lower bound. Numerical experiments on both synthetic and real-world datasets proved the effectiveness of the proposed algorithms and corroborated the theoretical analysis. PF-MAB is a new bandit framework that introduces the fundamental twist between local and global learning on top of the classical exploration-exploitation tradeoff, and sets the stage for potential future research activities.

## Acknowledgement

The CSs acknowledge the funding support by the US National Science Foundation under Grant ECCS-2029978 and ECCS-2033671. The work of JY was supported by the US National Science Foundation under Grant CNS-1956276 and ECCS-2030026.

## References

- Agarwal, A., Langford, J., and Wei, C.-Y. (2020). Federated residual learning. *arXiv preprint arXiv:2003.12880*.
- Bistriz, I. and Leshem, A. (2018). Distributed multiplayer bandits—a game of thrones approach. In *Advances in Neural Information Processing Systems*, pages 7222–7232.
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konecny, J., Mazzocchi, S., McMahan, H. B., Overveldt, T. V., Petrou, D., Ramage, D., and Roselander, J. (2019). Towards federated learning at scale: System design. In *Proceedings of the 2nd SysML Conference*, pages 1–15.
- Boursier, E., Kaufmann, E., Mehrabian, A., and Perchet, V. (2020). A practical algorithm for multiplayer bandits when arm means vary among players. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, Palermo, Sicily, Italy.
- Boursier, E. and Perchet, V. (2019). SIC-MMAB: synchronisation involves communication in multiplayer multi-armed bandits. In *Advances in Neural Information Processing Systems*, pages 12071–12080.
- Brânzei, S. and Peres, Y. (2019). Multiplayer bandit learning, from competition to cooperation. *arXiv preprint arXiv:1908.01135*.
- Cantador, I., Brusilovsky, P., and Kuflik, T. (2011). 2nd Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec 2011). In *Proceedings of the 5th ACM Conference on Recommender Systems*, RecSys 2011, New York, NY, USA. ACM.
- Deng, Y., Kamani, M. M., and Mahdavi, M. (2020). Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*.
- Dubey, A. and Pentland, A. (2020). Differentially-private federated linear bandits. *Advances in Neural Information Processing Systems*, 33.
- Fallah, A., Mokhtari, A., and Ozdaglar, A. (2020). Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948*.
- Hanzely, F. and Richtárik, P. (2020). Federated learning of a mixture of global and local models. *arXiv preprint arXiv:2002.05516*.
- Jiang, Y., Konečný, J., Rush, K., and Kannan, S. (2019). Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2019). Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*.
- Kulkarni, V., Kulkarni, M., and Pant, A. (2020). Survey of personalization techniques for federated learning. *arXiv preprint arXiv:2003.08673*.
- Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.*, 6(1):4–22.
- Landgren, P., Srivastava, V., and Leonard, N. E. (2016). On distributed cooperative decision-making in multiarmed bandits. In *2016 European Control Conference (ECC)*, pages 243–248. IEEE.
- Landgren, P., Srivastava, V., and Leonard, N. E. (2018). Social imitation in cooperative multiarmed bandits: partition-based algorithms with strictly local information. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 5239–5244. IEEE.
- Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60.
- Li, T., Song, L., and Fragouli, C. (2020). Federated recommendation system via differential privacy. In *IEEE International Symposium on Information Theory (ISIT)*, pages 2592–2597.
- Liu, K. and Zhao, Q. (2010). Distributed learning in multi-armed bandit with multiple players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681.
- Mansour, Y., Mohri, M., Ro, J., and Suresh, A. T. (2020). Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*.
- 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)*, pages 1273–1282, Fort Lauderdale, FL, USA.

- Rosenski, J., Shamir, O., and Szlak, L. (2016). Multi-player bandits—a musical chairs approach. In *International Conference on Machine Learning*, pages 155–163.
- Shahrampour, S., Rakhlin, A., and Jadbabaie, A. (2017). Multi-armed bandits in multi-agent networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2786–2790. IEEE.
- Shi, C. and Shen, C. (2021). Federated multi-armed bandits. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*.
- Shi, C., Xiong, W., Shen, C., and Yang, J. (2020). Decentralized multi-player multi-armed bandits with no collision information. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, Palermo, Sicily, Italy.
- Smith, V., Chiang, C.-K., Sanjabi, M., and Talwalkar, A. S. (2017). Federated multi-task learning. In *Advances in Neural Information Processing Systems*, pages 4424–4434.
- Wang, K., Mathews, R., Kiddon, C., Eichner, H., Beaufays, F., and Ramage, D. (2019). Federated evaluation of on-device personalization. *arXiv preprint arXiv:1910.10252*.
- Wang, Y., Hu, J., Chen, X., and Wang, L. (2020). Distributed bandit learning: Near-optimal regret with efficient communication. In *2020 International Conference on Learning Representations*.
- Zhu, Z., Zhu, J., Liu, J., and Liu, Y. (2020). Federated bandit: A gossiping approach. *arXiv preprint, arXiv:2010.12763*.