FCOM: A Federated Collaborative Online Monitoring Framework via Representation Learning

Tanapol Kosolwattana[1], Huazheng Wang[2], Raed Al Kontar[3], Ying Lin[1]
[1]Department of Industrial Engineering, University of Houston
[2]School of Electrical Engineering and Computer Science, Oregon State University
[3]Department of Industrial & Operations Engineering, University of Michigan
tkosolwattana@uh.edu, huazheng.wang@oregonstate.edu, alkontar@umich.edu, ylin53@central.uh.edu

Abstract

Monitoring a large population of dynamic processes with limited resources presents a significant challenge across various industrial sectors. This is due to 1) the inherent disparity between the available monitoring resources and the extensive number of processes to be monitored and 2) the unpredictable and heterogeneous dynamics inherent in the progression of these processes. Online learning approaches, commonly referred to as bandit methods, have demonstrated notable potential in addressing this issue by dynamically allocating resources and effectively balancing the exploitation of high-reward processes and the exploration of uncertain ones. However, most online learning algorithms are designed for 1) a centralized setting that requires data sharing across processes for accurate predictions or 2) a homogeneity assumption that estimates a single global model from decentralized data. To overcome these limitations and enable online learning in a heterogeneous population under a decentralized setting, we propose a federated collaborative online monitoring method. Our approach utilizes representation learning to capture the latent representative models within the population and introduces a novel federated collaborative UCB algorithm to estimate these models from sequentially observed decentralized data. This strategy facilitates informed monitoring of resource allocation. The efficacy of our method is demonstrated through theoretical analysis, simulation studies, and its application to decentralized cognitive degradation monitoring in Alzheimer's disease.

Code — https://github.com/TKosolwattana/FCOM_AAAI

Introduction

Monitoring a large population of dynamic processes within the constraints of monitoring resources poses a significant challenge across various industrial sectors, including healthcare and engineering systems (Lin, Liu, and Huang 2018; Kosolwattana, Wang, and Lin 2023). The complexity arises from two key factors: 1) the inherent disparity between the limited monitoring resources available and the large population of processes to be monitored, and 2) the uncertain and heterogeneous dynamics in the progression of these processes. In tackling this intricate problem, online learning from bandit feedback has demonstrated notable potential (Kosolwattana, Wang, and Lin 2023; Auer 2002). These methods offer a solution by dynamically allocating limited resources, effectively balancing the exploitation of processes yielding high rewards and the exploration of uncertain processes. While existing online learning algorithms, including those cited above, show promise, they still exhibit certain limitations that hinder their effectiveness in addressing the complexities of monitoring under resource constraints. The primary limitation of existing algorithms stems from their reliance on a centralized setting, necessitating the transmission of all data to a central server for model training. While this approach has demonstrated utility, it raises profound concerns about data privacy. In healthcare applications, such as patient health monitoring systems, this centralized model poses a substantial threat to data privacy, particularly with electronic health records (EHRs) containing sensitive patient information that ideally should remain within local organizations for privacy protection (Xu et al. 2021). Beyond privacy, issues related to the costs of communication and data storage in a central server further underscore the imperative for a decentralized approach (Kontar et al. 2021). Therefore, it is crucial to develop decentralized algorithms that empower the storage and analysis of monitoring data at the local level, mitigating privacy risks and reducing communication and storage costs.

An ostensibly straightforward method to achieve decentralized monitoring is by modeling and monitoring the progression dynamic of each process independently using conventional online learning algorithms (Wang et al. 2019). However, this approach can be less effective due to the uneven data distribution among individual processes, where the disparity in available data for each process can lead to suboptimal model performance. (Penny and Atkinson 2012). Secondly, training the model with stored data is less suitable for various real-world applications where data are usually available in an online manner with dynamic distributions (Zhang et al. 2021; Li et al. 2020). Given that each process features a heterogeneous dynamic that is not entirely independent, neglecting to exploit information sharing between processes may overlook opportunities to effectively span the model space for these processes' models (Xu et al. 2021). Failing to capitalize on inter-process information can result in a limited understanding of the broader context and interconnected dynamics within the monitored processes.

Several federated online learning algorithms have been developed in the literature to learn a global model under distributed datasets online. A notable approach in this context is federated multi-armed bandits (FMAB) (Shi, Shen, and Yang 2021; Zhu et al. 2021; Dubey and Pentland 2020), which provides a framework for units or processes to collaboratively solve a global bandit problem while retaining their information locally. In this approach, at each time step, each unit estimates the model locally and sends updates of model statistics to the central server. Subsequently, the central server sends back the aggregated statistics to enable each unit to enhance its model estimation for the next trial. This collaborative process allows the central server to develop a better selection strategy, striking a balance between exploration and exploitation. Nevertheless, these federated online learning algorithms usually focus on a homogeneous population that can be well described by a single global model. They can lead to a bias in estimating personalized models since the resulting global model only reflects the average effect, but fails to capture the between-unit variations (Tan et al. 2022). Some existing studies adopt a personalized modeling approach, such as the mixed model (Shi, Shen, and Yang 2021; Smith et al. 2017), which uses random effects to describe personalized variation. However, the mixed model used in these studies still cannot capture the underlying group structure in a heterogeneous population (Di et al. 2023). The underlying group structure represents diverse population characteristics where each group structure corresponds to one behavioral pattern or mechanism. Numerous studies have emphasized that capturing the underlying group structure could potentially enhance the accuracy of personalized modeling (Lin, Liu, and Huang 2018; Kumar, Yao, and Chu 2013).

To address this problem, we utilize the representation learning approach to exploit and monitor the hidden structure in a population. Representation learning is an approach that learns a representation to capture the common structure across different but related units (Du et al. 2020; Wang, Wu, and Wang 2017, 2016). It offers substantial benefits in learning individualized models from insufficient monitoring data, especially at the beginning of the monitoring period. However, to the best of our knowledge, there are no federated online learning algorithms that consider representation learning during unit modeling. Moreover, the units make their decisions locally and independently in existing federated online learning algorithms. Instead, the resource allocation problem requires decision-making under global constraint, i.e., monitoring dependent processes with limited monitoring resources, which differs from existing literature.

To mitigate these gaps, we propose a federated collaborative online monitoring framework. The proposed framework aims to allocate limited monitoring resources to a heterogeneous population in real-time, using federated online representation learning. It consists of three main phases: the representation learning-based modeling, the decentralized parameter estimation and communication, and the monitoring strategy design. In the first phase, we leverage a representation learning (Lin et al. 2018) approach to model the heterogeneous population, which assumes the existence of a shared representation that captures the underlying group structure among units. The second phase introduces a novel federated collaborative online monitoring (FCOM) algorithm to estimate the shared representation from distributed and sequentially arrived data. In the third phase, a novel upper confidence bound (UCB)-based score is developed to real-time assess the uncertainty of model estimation from the FCOM algorithm and inform the optimal monitoring resource allocation that balances the exploitation of high-reward units and exploration of uncertain ones. We rigorously prove sublinear regret and communication cost upper bounds of the proposed FCOM algorithm. The efficiency of the proposed method is validated by a simulation study and an empirical study on online cognitive degradation monitoring for Alzheimer's disease (AD).

The contributions of this paper include:

- Proposing a novel FCOM algorithm for decentralized online learning of latent group structures from distributed data, allowing local parameter updates and transmitting statistics that do not contain the true observed data to the central server.
- Developing an event-triggered communication strategy to reduce communication costs while maintaining low regret by allowing units to synchronize only when new information is sufficiently gathered.
- Demonstrating reduced upper regret bound compared to benchmark models through regret analysis when the latent group structure is low-rank.
- Validating method's effectiveness through simulation studies and an empirical study on adaptive cognitive monitoring in Alzheimer's disease, achieving comparable performance to benchmark models while preserving data privacy within local units.

Related Work

Federated Representation Learning

Federated representation learning is a learning framework that constructs a global model or shared representation from the aggregation of local units' updates to reflect common structures within a heterogeneous population. In cluster federated learning, the common representation is shown as the cluster of various global models where each global model reflects the unit's diverse pattern (Long et al. 2023). While this approach signifies group information and provides interpretability of population characteristics, it does not explicitly consider individual variations since each unit is only assigned to one global model. In multi-task federated learning, the central server broadcasts the information related to units to guide them in training their individualized models. For instance, (Yue, Kontar, and Gómez 2024) and (Smith et al. 2017) utilize a covariance matrix as a shared representation to provide similarity information between units. However, the estimated parameters from such approaches do not provide insights into the group structure to which each unit belongs. The most similar work to our study in this paper is the Population Personalized Federated Learning algorithm or PPFL in short (Di et al. 2023). The shared representation is constructed as a set of canonical models that demonstrate units' underlying mechanisms. Each unit has a membership vector, which measures the degree of the unit belonging to each canonical model. Even though these existing algorithms consider representation learning in various ways, they operate in offline settings, which is fundamentally different from our proposed algorithm since we are solving the problem of monitoring a subset of units in real-time.

Federated Online Learning

In federated online learning, units sequentially collect observations, update their local models, and send updates to a central server, which broadcasts a globally aggregated model at each time step (Mitra, Hassani, and Pappas 2021). Federated multi-armed bandits (FMAB) operate in a decentralized setting to preserve data privacy (Shi, Shen, and Yang 2021; Zhu et al. 2021). DisLinUCB in (Wang et al. 2019) considers a star-shaped communication to update globally shared parameters under synchronous setting, but they assume all units are homogeneous. Some existing algorithms assume heterogeneity among units by considering an additive model or a mixed model approach to account for both individual models and globally shared models (Li and Song 2022; Shi, Shen, and Yang 2021; Smith et al. 2017) with extension to kernel setting (Li et al. 2022, 2023). In Sync-LinUCB (Li and Wang 2022), the unknown parameter for each unit consists of a globally shared component, known as a fixed effect, and an individualized local component, known as a random effect. The random effect can be estimated from local data and the fixed effect needs to be estimated collaboratively in the central server. However, none of these existing algorithms can capture the inherent group structures of heterogeneous populations. Some studies incorporate clustering of bandits which group units with shared characteristics into clusters to handle the population heterogeneity (Blaser, Li, and Wang 2024; Yang et al. 2024; Liu et al. 2022; Korda, Szorenyi, and Li 2016). However, they assign each unit to a single cluster at each time step. Thus, each unit's coefficient is estimated using information from only one cluster, which does not explicitly reflect the underlying structure.

Method

Problem Statement

Consider a monitoring system featuring N units along with a central server responsible for allocating the monitoring resources over units and coordinating communication between itself and the units. Each unit's monitoring reward is a dynamic process associated with a time-varying feature vector $x_{it} \in \mathbb{R}^p$. If the unit is monitored, a corresponding monitoring reward can be observed from the environment, which is predicted by the feature vector using a personalized linear reward function $y_{it} = f_i(x_{it}) + \epsilon_{it}$. ϵ_{it} represents a random noise that follows a normal distribution. The personalized reward functions of N units are assumed to be related, meaning there exists a shared representation that captures the latent structure across the unit population. Exploiting this common representation can potentially improve the quality of each unit's model prediction, which leads to better monitoring decisions. The primary objective of the monitoring system is to select a set $\mathscr{A}_t$ in each trial t, comprising M units to be monitored, resulting in the maximal monitoring reward. The total monitoring reward can be represented as:

$$r_{\mathscr{A}_t} = \textstyle\sum_{i=1}^{N} a_{it} y_{it} \tag{1}$$

where a decision variable $a_{it} = 1$ if a unit $i \in \mathscr{A}_t$ or 0 otherwise and $\sum_{i=1}^{N} a_{it} = M$. Ideally, with the knowledge of reward functions, the optimal strategy of the central server is to choose the top M units with maximal rewards in each trial, denoted as a^*_{it}. However, f_i is unknown, and the central server needs to update its selection strategy through the learning of models from observation (x_{it}, y_{it}) in each trial t. Thus, the goal of the central server is equivalent to minimizing the cumulative regret $R(T)$ which is defined as the difference between the central server's total reward and the total reward of the optimal strategy shown as follows.

$$R(T) = \textstyle\sum_{t=1}^{T} R_t = \sum_{t=1}^{T} \sum_{i=1}^{N} a^*_{it} y_{it} - a_{it} y_{it} \tag{2}$$

To address data privacy concerns between units and the central server, the monitoring system ensures that historical information, including the unit's observed feature vectors and rewards, remains confidential and distributed within each unit, not exposed to the central server or other units. The online monitoring problem becomes more complex under distributed data due to the need for efficient communication between local units and the central server and increased uncertainty in model estimation. To solve this problem, this paper presents a federated collaborative online monitoring approach that integrates representation learning and a federated online learning algorithm.

Federated Collaborative Online Monitoring

Representation Learning-Based Modeling This paper assumes the reward function of each unit is linear with respect to the feature vectors i.e., $f_i(x_{it}) = \beta_i^T x_{it}$ where $\beta_i \in \mathbb{R}^p$ is the coefficient vector of unit i. As the units are related to each other, extracting the common structure between them could help their corresponding models learn more efficiently than considering each unit independently (Yang et al. 2021). Following the representation learning approach introduced in (Yang et al. 2021), β_i is denoted as $\beta_i = Qc_i$ where $Q = [q^{(1)}, \ldots, q^{(K)}] \in \mathbb{R}^{p \times K}$ is a low-rank matrix that represents K representative models ($K \ll p$) in the reward functions and c_i serves as the unit i's coefficient. This representation captures the underlying group structure among units, a subset of which may follow the same representative model. Each unit's coefficients on K representative models are denoted as a membership vector $\{c_i\}_{i=1}^N, c_i \in \mathbb{R}^K$. Thus, the unit's reward function can be illustrated as the weight combination of K representative models $f_i(x_{it}) = \sum_k c_{ik} x_{it}^T q^{(k)}$. In the monitoring problem, we aim to learn the representative models and units' membership vectors to make optimal decisions. Assuming parameters in K representative models, Q, and units' membership vectors, $\{c_i\}_{i=1}^N$, are unconstrained, such a problem can be formulated as a global non-convex optimization problem in the following

$$\min_{C,Q} \textstyle\sum_t \sum_i \|x_{it} Q c_i - y_{it}\|_2^2 + \eta_1 \|Q\|_F^2 + \eta_2 \|C\|_F^2 \tag{3}$$

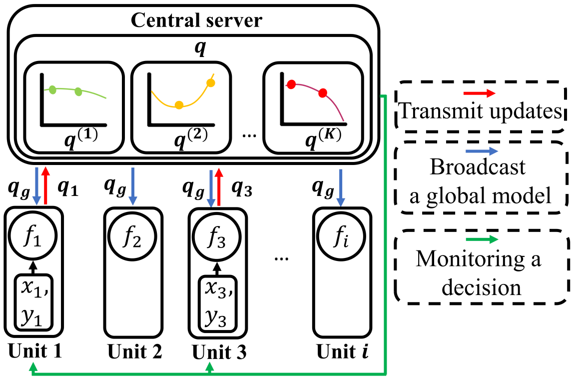

Figure 1: The overall framework of the proposed Federated Collaborative Online Monitoring (FCOM) algorithm

where $C = [c_1, \ldots, c_N]$. The L2 regularization is used on Q and C to control the scale of unknown parameters with tuning parameters η_1 and η_2.

Few federated learning algorithms are developed to solve the optimization problem similar to Equation 3 from distributed data (Di et al. 2023; Liang et al. 2020; Jeong and Hwang 2022). Generally, they alternatively update the shared representation and membership vectors in each step. Equation 3 can be achieved by decomposing into a set of sub-problems that can be solved using local unit data. The feature matrix $\mathbf{X}_{it} = [0, \ldots, x_{it}, \ldots, 0] \in \mathbb{R}^{p \times N}$ is transformed to $X_{it} = \mathbf{X}_{it} \otimes I_K \in \mathbb{R}^{Kp \times NK}$ where I_K is the $K \times K$ identity matrix. Let $\tilde{c}_i = [0, \ldots, c_i, \ldots, 0] \in \mathbb{R}^{NK \times 1}$, and $q = vec(Q) = [q^{(1)T}, \cdots, q^{(K)T}]^T \in \mathbb{R}^{Kp \times 1}$ where $vec(\cdot)$ is the vectorized operation, the sub-problem for each local unit i is represented as follows.

$$\min_{\tilde{c}_i, q_i} \sum_{t=1}^{T} \left\| \tilde{c}_i^T X_{it}^T q_i - y_{it} \right\|_2^2 + \eta_1 \|q_i\|_2^2 + \eta_2 \|c_i\|_2^2 \quad (4)$$

However, these algorithms solve the representation learning problem in an offline setting, which is inappropriate for online monitoring where the models need to be updated with sequentially arrived data and decisions need to be made under exploitation-exploration trade-off. Thus, we propose a novel federated collaborative online monitoring algorithm to solve this issue.

FCOM Algorithm We propose a federated collaborative online monitoring (FCOM) algorithm to estimate q that is shared across N units and c_i which is specified for each unit from the distributed and sequentially observed data, as shown in Fig. 1. Assuming that a subset of units $\mathscr{A}_t$ is monitored in trial t with new observations of monitoring rewards and feature vectors, the proposed method initially updates c_i for monitored units and their local estimations of q, denoted as q_i, using an alternative least squares (ALS) algorithm that provides closed-form solutions as demonstrated in **Step-1**. Then, in **Step-2**, the updated estimations of q_i in **Step-1** which are potential to benefit the global estimation are further uploaded to the central server. The central server uses these updates to estimate and broadcast the global representation q_g for all units so both monitored and not-monitored units collaboratively update models via new observations.

Step-1 Local update: There are two steps to estimate parameters in representation learning using an ALS algorithm. It firstly estimates $\tilde{c}_{it}$ by fixing q_{it} in **Step-1.1** and then estimates q_{it} with fixed $\tilde{c}_{it}$ in **Step-1.2**.

Step-1.1: When a local estimation of the representation q_{it} is fixed, the estimation of a unit i's membership vector $\tilde{c}_{it}$ can be obtained by solving the objective function in Equation 4. The closed-form solution of $\hat{\tilde{c}}_{it}$ can be achieved by $\hat{\tilde{c}}_{it} = D_{it}^{-1} d_{it}$ in which,

$$D_{it} = \textstyle\sum_{t'=1}^{t} X_{it'}^T q_{it} q_{it}^T X_{it'} + \eta_2 I_{NK} \quad (5)$$

$$d_{it} = \textstyle\sum_{t'=1}^{t} X_{it'}^T q_{it} y_{it'} \quad (6)$$

Step-1.2: When a unit i's membership vector $\tilde{c}_{it}$ is fixed, the estimation of the local representation q_{it} can be obtained by solving the objective function in Equation 4. The closed-form solution of $\hat{q}_{it}$ can be achieved by $\hat{q}_{it} = A_{it}^{-1} b_{it}$ in which,

$$A_{it} = \textstyle\sum_{t'=1}^{t} X_{it'} \tilde{c}_{it} \tilde{c}_{it}^T X_{it'}^T + \eta_1 I_{Kp} \quad (7)$$

$$b_{it} = \textstyle\sum_{t'=1}^{t} X_{it'} \tilde{c}_{it} y_{it'} \quad (8)$$

The unit i alternatively solves these two steps until reaching convergence. The procedure of **Step-1** is summarized in Line 11 - Line 15 of **Algorithm 1**.

Step-2 Global update: The representation q is shared across all N units and needs to be estimated collaboratively. The update $\{\Delta A_{it}, \Delta b_{it}\}$ is defined as the information of q that is learned from the local estimation of $\hat{q}_{it}$ and denoted as

$$\Delta A_{it} = \Delta A_{it-1} + \textstyle\sum_{t'=1}^{t} X_{it'} \tilde{c}_{it} \tilde{c}_{it}^T X_{it'}^T \quad (9)$$

$$\Delta b_{it} = \Delta b_{it-1} + \textstyle\sum_{t'=1}^{t} X_{it'} \tilde{c}_{it} y_{it'} \quad (10)$$

To avoid uploading new information every trial and reduce the communication cost, we introduce an event trigger for communication: only communicate/synchronize when local units have gathered sufficient new information. The evaluation is in the form of the matrix determinant-based criterion, as shown in Equation 11 where $\gamma_{U_q} \geq 1$ is a hyperparameter:

$$\det(A_{it}) > \gamma_{U_q} \det(A_{it} - \Delta A_{it}) \quad (11)$$

If the unit i passes Equation 11, it uploads $\{\Delta A_{it}, \Delta b_{it}\}$ to the central server and reset $\Delta A_{it} = 0, \Delta b_{it} = 0$. The central server aggregates $\{\Delta A_{it}, \Delta b_{it}\}$ to its corresponding statistics $\{A_{gt}, b_{gt}\}$. After the aforementioned procedures are repeated to all monitored units, the central server estimates the global representation $\hat{q}_{gt} = A_{gt}^{-1} b_{gt}$ where $A_{gt} \mathrel{+}= \Delta A_{it}, b_{gt} \mathrel{+}= \Delta b_{it}$. Finally, the central server broadcasts $\hat{q}_{gt}$ and its corresponding statistics to all units. Every unit $j \in [N]$ updates $\hat{q}_{jt+1} = \hat{q}_t, A_{jt+1} = A_{gt}, b_{jt+1} = b_{gt}$. The procedure of **Step-2** is summarized in Line 16 - Line 24 of **Algorithm 1**.

Remark: Algorithm 1 offers more flexibility than algorithms operating in a fully synchronous setting (Wang et al. 2019; Dubey and Pentland 2020). In our approach, all units are not required to upload their statistics whenever one unit triggers communication. This allows units with less informative statistics to refrain from uploading, thereby enhancing the efficiency of communication.

Algorithm 1: Federated collaborative online monitoring (FCOM) algorithm

1: **Input:** $\eta_1, \eta_2, \gamma_{U_q}, \alpha^q, \alpha^{\tilde{c}}, M$
2: **Initialization:** $A_{i0} \leftarrow \eta_1 I_{Kp}$, $b_{i0} \leftarrow 0$, $D_{i0} \leftarrow \eta_2 I_{NK}$, $d_{i0} \leftarrow 0, \forall i \in [N]$; $\Delta A_{i0} \leftarrow 0$, $\Delta b_{i0} \leftarrow 0$, $\Delta D_{i0} \leftarrow 0$, $\Delta d_{i0} \leftarrow 0, \forall i \in [N]$; $A_{g0} \leftarrow \eta_1 I_{Kp}, b_{g0} \leftarrow 0$
3: **for** $t \leftarrow 1, \dots, T$ **do**
4:   **for** all $i \in [N]$ **do**
5:     Transform x_{it} to $X_{it} \in \mathbb{R}^{Kp \times NK}$
6:     Compute $\hat{u}_{it} \leftarrow \hat{\tilde{c}}_{it}^T X_{it}^T \hat{q}_{it} + \alpha^{\tilde{c}} \sqrt{\hat{q}_{it}^T X_{it} D_{it}^{-1} X_{it}^T \hat{q}_{it}} + \alpha^q \sqrt{\hat{\tilde{c}}_{it}^T X_{it}^T A_{it}^{-1} X_{it} \hat{\tilde{c}}_{it}}$
7:   **end for**
8:   The central server selects the subset $\mathscr{A}_t$ that includes M units with highest UCB score computed in Line 6.
9:   Each selected unit i' receives $y_{i't}$ from the environment.
10:   **for** all $i' \in \mathscr{A}_t$ **do**
11:     **while** not converge **do**
12:       Update $D_{i't}, d_{i't}, A_{i't}, b_{i't}$ by Equation 5, 6 7, 8
13:       Update $\Delta A_{i't}, \Delta b_{i't}$ using Equation 9, 10
14:       Compute $\hat{q}_{i't} \leftarrow A_{i't}^{-1} b_{i't}, \hat{\tilde{c}}_{i't} \leftarrow D_{i't}^{-1} d_{i't}$
15:     **end while**
16:     **if** $\det(A_{i't}) > \gamma_{U_q} \det(A_{i't} - \Delta A_{i't})$ **then**
17:       $A_{gt} \mathrel{+}= \Delta A_{i't}, b_{gt} \mathrel{+}= \Delta b_{i't}$
18:       Reset $\Delta A_{i't} \leftarrow 0, \Delta b_{i't} \leftarrow 0$
19:     **end if**
20:   **end for**
21:   **if** at least one unit triggers line 16 **then**
22:     The central server estimates $\hat{q}_{gt} \leftarrow A_{gt}^{-1} b_{gt}$ and sends $\{A_{gt}, b_{gt}, \hat{q}_{gt}\}$ to all units $i \in [N]$.
23:     All units $i \in [N]$ update $\hat{q}_{it+1}$ and $\{A_{it+1}, b_{it+1}\}$.
24:   **end if**
25: **end for**

Monitoring Strategy Design Based on the estimated representation and membership vectors, each unit can obtain a prediction of the monitoring reward in the next trial conditioning on its feature vector, i.e. $\hat{y}_{it+1} = \hat{c}_{it+1}^T X_{it+1}^T \hat{q}_{it+1}$. However, collecting only predicted rewards does not guarantee optimal selection, as it neglects prediction uncertainty estimated by each local unit. To address this issue, the upper confidence bound (UCB) principle is commonly used, which comprises two main components: the unit's predicted reward and the unit's prediction uncertainty (Auer 2002). This score balances the objective of selecting units with the highest rewards (exploitation) and units with high uncertainty in reward predictions (exploration). To obtain the UCB score, we derive the uncertainty of parameters estimated from distributed data using the algorithm in Section . The uncertainty of the estimated parameters is composed of two parts: $\|\hat{q}_{it} - q^*\|_{A_{it}}$ and $\|\hat{\tilde{c}}_{it} - \tilde{c}^*\|_{D_{it}}$ where q^* and $\tilde{c}^*$ are the ground-truth parameters. Based on the closed-form solutions in the ALS algorithm, the confidence ellipsoid of $\hat{q}_{it}$ and $\hat{\tilde{c}}_{it}$ can be obtained by **Lemma 1** below.

Lemma 1 *When the Hessian matrices of the objective function in Equations 4 are positive definite at the optimizer q^* and $\tilde{c^*}$, for any $\epsilon_1 \geq 0$, $\epsilon_2 \geq 0$, $0 < v_1 < 1$, $0 < v_2 < 1$, $\|X_t\|_2 \leq S$, $\|q_{it}\|_2 \leq L$, $\|\tilde{c}_{it}\|_2 \leq P$, and for any $\delta \geq 0$, with probability at least 1 - δ, the estimation errors of $\hat{q}_{it}$ and $\hat{\tilde{c}}_{it}$ are upper bounded by:*

$$\|\hat{q}_{it} - q^*\|_{A_{it}} \leq \sqrt{Kp \ln(\frac{\eta_1 Kp + tS^2P^2}{\eta_1 Kp\delta})} + \sqrt{\eta_1} L + \frac{2SPL}{\sqrt{\eta_1}} \frac{(v_1+\epsilon_1)(1-(v_1+\epsilon_1)^t)}{1-(v_1+\epsilon_1)}$$

$$\|\hat{\tilde{c}}_{it} - \tilde{c}^*\|_{D_{it}} \leq \sqrt{NK \ln(\frac{\eta_2 NK + tS^2L^2}{\eta_2 NK\delta})} + \sqrt{\eta_2} P + \frac{2SPL}{\sqrt{\eta_2}} \frac{(v_2+\epsilon_2)(1-(v_2+\epsilon_2)^t)}{1-(v_2+\epsilon_2)}$$

where α^q and $\alpha^{\tilde{c}}$ are tuning parameters that are defined as the upper bounds of $\|\hat{q}_{it} - q^*\|_{A_{it}}$ and $\|\hat{\tilde{c}}_{it} - \tilde{c}^*\|_{D_{it}}$ respectively. For the monitoring strategy, the central server selects M units that have the highest UCB score. The unit i's UCB score is defined as follows:

$$\hat{u}_{it} = \hat{\tilde{c}}_{it}^T X_{it}^T \hat{q}_{it} + \alpha^{\tilde{c}} \sqrt{\hat{q}_{it}^T X_{it} D_{it}^{-1} X_{it}^T \hat{q}_{it}} + \alpha^q \sqrt{\hat{\tilde{c}}_{it}^T X_{it}^T A_{it}^{-1} X_{it} \hat{\tilde{c}}_{it}} \quad (12)$$

Equation 12 illustrates the UCB score $\hat{u}_{it}$, which consists of three parts: the predicted reward of unit i, the prediction uncertainty due to the unit's membership vector, and the prediction uncertainty due to representation estimations. Since the resulting UCB score for each unit calculated from the estimated parameters is only exposed to the central server, it ensures privacy among all units as the central server cannot directly access the raw information (i.e., the unit's feature matrix and the unit's actual reward). The procedure of FCOM algorithm is summarized in **Algorithm 1**.

Regret Analysis

Theorem 1 *For any $\delta > 0$, with the probability at least 1 - δ, the cumulative regret and the communication cost of the FCOM algorithm is upper bounded by:*

$$R(T) \leq 2\alpha^q G \sqrt{2MTKp \ln(1 + \frac{TS^2P^2}{\eta_1 Kp})} + 2\alpha^{\tilde{c}} \sqrt{2MTNK \ln(1 + \frac{TS^2L^2}{\eta_2 NK})} + \frac{2Mv^2(1-v^{2T})}{1-v^2}$$

$$C(T) \leq \lceil \frac{N}{M} \rceil \frac{(N+1)Kp}{\ln(\gamma_{U_q})} \ln(1 + TS^2P^2/(\eta_1 Kp))$$

where $G = \sqrt{1 + (M-1)(\gamma_{U_q} - 1)}$.

Theorem 1 presents the upper regret bound and communication cost of the FCOM algorithm, providing a theoretical analysis of the proposed selection strategy's effectiveness. The upper regret bound is decomposed into three components. The first term reflects the regret stemming from the estimation uncertainty of shared representation q, leveraging communication between units and the central server. The second term accounts for the regret associated with the estimation uncertainty of units' membership vectors $\tilde{c}$. The third term is the regret due to q-linear convergence in estimating $\tilde{c}$ and q over T trials. Compared to the existing federated online learning algorithm, Sync-LinUCB, which utilizes a mixed model for reward estimation and achieves an upper regret bound of $O(Np\sqrt{T}\ln T)$ (Li and Wang 2022), the FCOM algorithm offers an improved regret bound of $O(NK\sqrt{T}\ln T)$ by capturing the shared representation among units. Additionally, it maintains a comparable order of communication cost at $O(N^3 p \ln T)$ using $\gamma_{U_q} = e^{K/N}$. This improvement is notable under the assumption that $K \ll p$. However, if the latent group structure is not low-rank, FCOM may result in regret bounds similar to Sync-LinUCB.

Experiments

We evaluate the proposed method using both a simulation study and a real-world case study of decentralized cognitive degradation monitoring of Alzheimer's Disease (AD) in older populations. In both studies, we consider 100 local devices corresponding to individual units and a central server. The monitoring process begins with each local device sending the UCB score of its respective unit to the central server. The central server then selects the top M local devices with the highest UCB scores and instructs them to collect measurements and estimate their models locally. Each chosen device checks its communication conditions and, if satisfactory, sends its model update to the central server. The central server aggregates all received updates to estimate the global representation model, which is then broadcast to all units. This process is repeated until all monitoring trials are complete. We explore two monitoring scenarios to reflect varying levels of resource allocation in the real world. In the high-constraint scenario, only 33% of the total population is monitored in each monitoring trial ($M = 33\%$). In contrast, the relaxed-constraint scenario allows more units to be monitored, encompassing 66% of the total population ($M = 66\%$). We repeatedly monitor the populations over 30000 trials ($T = 30000$) in all experiments. The proposed FCOM algorithm is compared against various benchmark MAB algorithms, including LinUCB (Li et al. 2010), Sync-LinUCB (Li and Wang 2022), and CLUCB (Kosolwattana, Wang, and Lin 2023). LinUCB estimates the coefficients of units' reward functions independently through L2-regularized least square estimation. The Sync-LinUCB is also a decentralized online modeling and monitoring algorithm that models personalized reward functions using a mixed model. Specifically, random effects are estimated within local units independently whereas the fixed effect is estimated in the central server globally. However, neither algorithm exploits the representation learning approach to extract the latent patterns presented in the unit's reward function. CLUCB considers online modeling and monitoring of units' reward functions via representation learning. However, it is developed for the centralized setting, meaning it shares and saves all units' observations of reward and feature vectors in the central server. The cumulative regret defined in Equation 2 is used to compare the performance of all algorithms in both simulation and real-world studies. The communication cost, defined as the number of times updates are transferred between units and the central server, is also compared to show the effectiveness of implementing a federated learning approach to solve the monitoring problem.

Simulation Studies

We build the reward functions of 100 units ($N = 100$) and simulate their observations of monitoring reward over time, which is represented as $y_{it} = x_{it}\beta_i + \epsilon_{it}$. The feature vector $x_{it} = [x_{i1t}, x_{i2t}, \ldots, x_{ipt}]$ is simulated from p sigmoid functions which are formulated by (Young et al. 2015; Venkatraghavan et al. 2019) $x_{ijt} = a_j + \frac{r_j}{1+\exp(-d_j(t-c_j))} + \epsilon_{it}$ where parameters a_j, r_j, d_j, c_j and noise ϵ_{it} are randomly generated from a standard normal distribution. To guarantee the reward functions share a latent group structure, the coefficients in reward functions are assumed to be generated from K representative models and a membership vector, denoted as $\beta_i = Qc_i$. We consider three representative models ($K = 3$) in this experiment. The representative models are randomly generated from the multivariate normal distribution. To ensure each unit's reward function can be effectively resembled by one of the representative models, its membership vector is randomly generated from a mixture of Gaussian distributions. Specifically, the Gaussian distributions have zero-mean components, and the respective prior probabilities are estimated from the real data: $F_1(c) \sim N(0, \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}), F_2(c) \sim N(0, \begin{bmatrix} 1 & 0 & 0 \\ 0 & \sigma^2 & 0 \\ 0 & 0 & 1 \end{bmatrix})$, $F_3(c) \sim N(0, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \sigma^2 \end{bmatrix})$. The covariance matrix in each distribution is a diagonal matrix with a k^{th} diagonal element equal to σ^2 and the other elements equal to 1. σ^2 controls the significance of the latent group structure, where a larger value indicates more distinct groups. In this experiment, we set σ^2 at 100 to simulate three distinct groups. Random noises are generated from a standard normal distribution. We utilize k-means clustering to assign the initial value of c_i based on the centroids learned from the similarity matrix where each element in the matrix is calculated by the cosine similarity of the membership vectors. We then use c_i to initialize the value of Q by following the update part from lines 12 and 14 in **Algorithm 1**. The experiment is repeated three times to estimate the variation in performance.

In Figure 2a) and b), the cumulative regret of various online modeling and monitoring algorithms is illustrated under two monitoring scenarios. Sync-LinUCB exhibits lower cumulative regret than LinUCB in both scenarios but shows high variations in the high-constraint scenario. Our proposed FCOM algorithm achieves a lower cumulative regret than LinUCB and Sync-LinUCB. Figure 2c) shows that the communication costs of the proposed method can be reduced by increasing the threshold. Moreover, the cumulative regret of our method is more stable and lower than Sync-LinUCB under various communication costs. The results indicate that incorporating shared representation and providing information about the group structure among units leads to the most

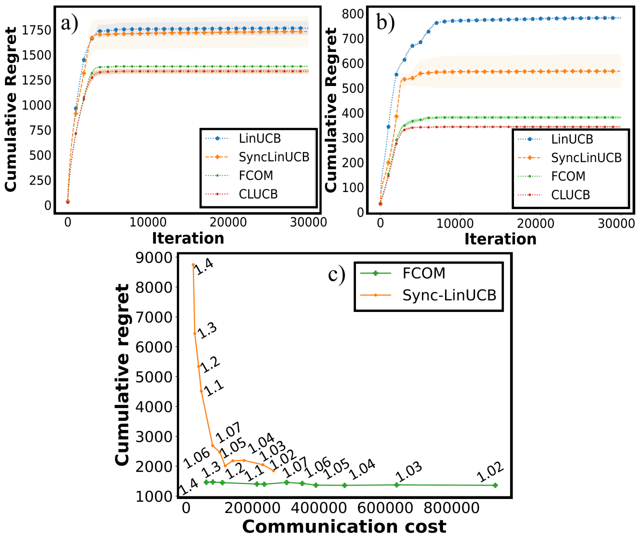

Figure 2: a)-b) The convergence of cumulative regret at $M = 33\%$ and $M = 66\%$, c) Communication cost and cumulative regret under various thresholds γ_{U_q} at $T = 30000$.

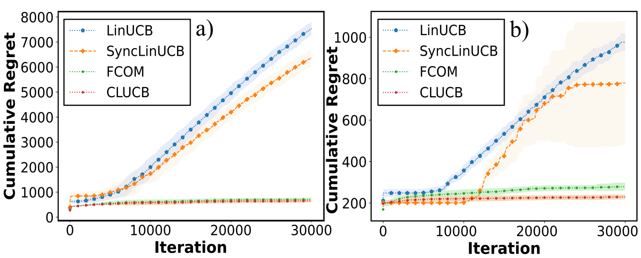

Figure 3: The cumulative regret for all algorithms shown in the AD dataset with a) $M = 33\%$ and b) $M = 66\%$

significant performance improvement. Moreover, compared to CLUCB which is in a centralized setting, the cumulative regret of FCOM is slightly higher. This suggests modeling the dynamic health progression in a federated learning setting to preserve data privacy does not hurt the performance of the proposed algorithm.

Application to Online Cognitive Monitoring in Alzheimer's Disease (AD)

A case study demonstrates the application of the proposed method for decentralized online monitoring of cognitive degradation in older adults. Cognitive monitoring via cognitive measurements, such as the Mini-Mental State Examination (MMSE), is essential for early detection of Alzheimer's disease (AD). However, existing cognitive monitoring relies on 1) routine hospital visits of all older adults; 2) a centralized setting that collects and stores all cognitive measurements in a central server of a hospital. To enable cognitive monitoring under limited clinical resources and enhance patients' privacy, the proposed method is applied, with the objective of allocating the monitoring resources to patients having severe cognitive degradation. The monitoring reward is defined based on the cognitive measurement (MMSE score), with a lower MMSE score corresponding to a more severe status (Mitchell 2009). A more severe cognitive status corresponds to a higher monitoring reward, facilitating early prevention and treatment of AD. The progression of MMSE is modeled using a degradation model with time and polynomial basis functions as predictors. Specifically, the MMSE measurement y_{it} of the patient i at time t is formulated as (Biesanz et al. 2004): $y_{it} = \beta_{i0} + \beta_{i1}t + \beta_{i2}t^2 + \beta_{i3}t^3 + \beta_{i4}t^4 + \beta_{i5}t^5 + \epsilon_{it}$. The dataset, acquired from the Alzheimer's Disease Neuroimaging Initiative (ADNI), includes longitudinal MMSE measurements for older adults. We sample 100 older adults from ADNI, with their MMSE measurements taken at baseline, $12^{th}, 24^{th}, 36^{th}, 48^{th}$, and 60^{th} month. There are distinct cognitive degradation patterns within the population due to three patient types: the normal aging patients (34%), the patients with mild cognitive impairment (33%), and the patients with Alzheimer's disease (33%) (Mitchell 2009). Therefore, the number of representative models is fixed at 3. The MMSE measurements collected from baseline to the 60^{th} month are equally divided into 30,000 trials. We use linear interpolation to impute the missing MMSE values for each older adult. To estimate performance variation, we repeat the experiments five times.

The cumulative regrets of different algorithms under two monitoring scenarios are compared in Figure 3. LinUCB and Sync-LinUCB exhibit higher cumulative regret than FCOM in both scenarios. The cumulative regret of Sync-LinUCB reaches convergence only in the high-constraint scenario, while our proposed algorithm converges and outperforms Sync-LinUCB in both scenarios. This observation suggests that effectively leveraging representation learning improves the accuracy of cognitive degradation modeling in AD and enables more accurate monitoring resource allocation. Moreover, it indicates the stability of the proposed algorithm regardless of the number of selected older adults. Moreover, similar to the simulation studies, the cumulative regret of FCOM is slightly higher than CLUCB, suggesting that modeling older adults' cognitive degradation in a decentralized setting to preserve data privacy does not hurt the performance of the proposed algorithm.

Conclusions

This paper introduces a novel Federated Online Collaborative Monitoring (FCOM) algorithm for decentralized online learning of a population with latent structured units. The FCOM algorithm uses representation learning to capture the latent group structure and estimates this structure from distributed, sequentially observed data by updating real-time local statistics and communicating these statistics between the central server and local units. A new UCB score is developed to optimize monitoring resource allocation. Regret analysis shows that our algorithm achieves a reduced upper regret bound compared to benchmark models when the latent group structure is low-rank. Simulation studies and an empirical study of online cognitive monitoring in Alzheimer's disease validate that our algorithm achieves lower cumulative regret compared to benchmark models.

Acknowledgements

Huazheng Wang is supported by National Science Foundation under grant IIS-2403401.

Raed Al Kontar is supported by National Science Foundation under grant 2144147 and 2328010.

Ying Lin is supported by the US Department of Transportation (USDOT) Tier-1 University Transportation Center (UTC) Transportation Cybersecurity Center for Advanced Research and Education (CYBER-CARE) (Grant No. 69A3552348332).

References

Auer, P. 2002. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov): 397–422.

Biesanz, J. C.; Deeb-Sossa, N.; Papadakis, A. A.; Bollen, K. A.; and Curran, P. J. 2004. The role of coding time in estimating and interpreting growth curve models. *Psychological methods*, 9(1): 30.

Blaser, E.; Li, C.; and Wang, H. 2024. Federated Linear Contextual Bandits with Heterogeneous Clients. In *International Conference on Artificial Intelligence and Statistics*, 631–639. PMLR.

Di, H.; Yang, Y.; Ye, H.; and Chang, X. 2023. PPFL: A Personalized Federated Learning Framework for Heterogeneous Population. *arXiv preprint arXiv:2310.14337*.

Du, S. S.; Hu, W.; Kakade, S. M.; Lee, J. D.; and Lei, Q. 2020. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*.

Dubey, A.; and Pentland, A. 2020. Differentially-private federated linear bandits. *Advances in Neural Information Processing Systems*, 33: 6003–6014.

Jeong, W.; and Hwang, S. J. 2022. Factorized-FL: Personalized Federated Learning with Parameter Factorization & Similarity Matching. *Advances in Neural Information Processing Systems*, 35: 35684–35695.

Kontar, R.; Shi, N.; Yue, X.; Chung, S.; Byon, E.; Chowdhury, M.; Jin, J.; Kontar, W.; Masoud, N.; Nouiehed, M.; et al. 2021. The internet of federated things (IoFT). *IEEE Access*, 9: 156071–156113.

Korda, N.; Szorenyi, B.; and Li, S. 2016. Distributed clustering of linear bandits in peer to peer networks. In *International conference on machine learning*, 1301–1309. PMLR.

Kosolwattana, T.; Wang, H.; and Lin, Y. 2023. Online Modeling and Monitoring of Dependent Processes under Resource Constraints. *arXiv preprint arXiv:2307.14208*.

Kumar, A.; Yao, W.; and Chu, C.-H. 2013. Flexible process compliance with semantic constraints using mixed-integer programming. *INFORMS Journal on Computing*, 25(3): 543–559.

Li, C.; and Wang, H. 2022. Asynchronous Upper Confidence Bound Algorithms for Federated Linear Bandits. In Camps-Valls, G.; Ruiz, F. J. R.; and Valera, I., eds., *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, 6529–6553. PMLR.

Li, C.; Wang, H.; Wang, M.; and Wang, H. 2022. Communication efficient distributed learning for kernelized contextual bandits. *Advances in Neural Information Processing Systems*, 35: 19773–19785.

Li, C.; Wang, H.; Wang, M.; and Wang, H. 2023. Learning kernelized contextual bandits in a distributed and asynchronous environment. In *International Conference on Learning Representation*.

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*, 661–670.

Li, L.; Fan, Y.; Tse, M.; and Lin, K.-Y. 2020. A review of applications in federated learning. *Computers & Industrial Engineering*, 149: 106854.

Li, T.; and Song, L. 2022. Privacy-preserving communication-efficient federated multi-armed bandits. *IEEE Journal on Selected Areas in Communications*, 40(3): 773–787.

Liang, P. P.; Liu, T.; Ziyin, L.; Allen, N. B.; Auerbach, R. P.; Brent, D.; Salakhutdinov, R.; and Morency, L.-P. 2020. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*.

Lin, Y.; Liu, K.; Byon, E.; Qian, X.; Liu, S.; and Huang, S. 2018. A Collaborative Learning Framework for Estimating Many Individualized Regression Models in a Heterogeneous Population. *IEEE Transactions on Reliability*, 67(1): 328–341.

Lin, Y.; Liu, S.; and Huang, S. 2018. Selective sensing of a heterogeneous population of units with dynamic health conditions. *IISE Transactions*, 50(12): 1076–1088.

Liu, X.; Zhao, H.; Yu, T.; Li, S.; and Lui, J. C. 2022. Federated online clustering of bandits. In *Uncertainty in Artificial Intelligence*, 1221–1231. PMLR.

Long, G.; Xie, M.; Shen, T.; Zhou, T.; Wang, X.; and Jiang, J. 2023. Multi-center federated learning: clients clustering for better personalization. *World Wide Web*, 26(1): 481–500.

Mitchell, A. J. 2009. A meta-analysis of the accuracy of the mini-mental state examination in the detection of dementia and mild cognitive impairment. *Journal of psychiatric research*, 43(4): 411–431.

Mitra, A.; Hassani, H.; and Pappas, G. J. 2021. Online federated learning. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 4083–4090. IEEE.

Penny, K. I.; and Atkinson, I. 2012. Approaches for dealing with missing data in health care studies. *Journal of clinical nursing*, 21(19pt20): 2722–2729.

Shi, C.; Shen, C.; and Yang, J. 2021. Federated Multi-armed Bandits with Personalization. arXiv:2102.13101.

Smith, V.; Chiang, C.-K.; Sanjabi, M.; and Talwalkar, A. S. 2017. Federated multi-task learning. *Advances in neural information processing systems*, 30.

Tan, A. Z.; Yu, H.; Cui, L.; and Yang, Q. 2022. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*.

Venkatraghavan, V.; Bron, E. E.; Niessen, W. J.; Klein, S.; Initiative, A. D. N.; et al. 2019. Disease progression timeline estimation for Alzheimer's disease using discriminative event based modeling. *NeuroImage*, 186: 518–532.

Wang, H.; Wu, Q.; and Wang, H. 2016. Learning hidden features for contextual bandits. In *Proceedings of the 25th ACM international on conference on information and knowledge management*, 1633–1642.

Wang, H.; Wu, Q.; and Wang, H. 2017. Factorization bandits for interactive recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Wang, Y.; Hu, J.; Chen, X.; and Wang, L. 2019. Distributed bandit learning: Near-optimal regret with efficient communication. *arXiv preprint arXiv:1904.06309*.

Xu, J.; Glicksberg, B. S.; Su, C.; Walker, P.; Bian, J.; and Wang, F. 2021. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5: 1–19.

Yang, H.; Liu, X.; Wang, Z.; Xie, H.; Lui, J. C.; Lian, D.; and Chen, E. 2024. Federated Contextual Cascading Bandits with Asynchronous Communication and Heterogeneous Users. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20596–20603.

Yang, J.; Hu, W.; Lee, J. D.; and Du, S. S. 2021. Impact of Representation Learning in Linear Bandits. arXiv:2010.06531.

Young, A. L.; Oxtoby, N. P.; Ourselin, S.; Schott, J. M.; Alexander, D. C.; Initiative, A. D. N.; et al. 2015. A simulation system for biomarker evolution in neurodegenerative disease. *Medical image analysis*, 26(1): 47–56.

Yue, X.; Kontar, R. A.; and Gómez, A. M. E. 2024. Federated data analytics: A study on linear models. *IISE Transactions*, 56(1): 16–28.

Zhang, C.; Xie, Y.; Bai, H.; Yu, B.; Li, W.; and Gao, Y. 2021. A survey on federated learning. *Knowledge-Based Systems*, 216: 106775.

Zhu, Z.; Zhu, J.; Liu, J.; and Liu, Y. 2021. Federated bandit: A gossiping approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 5(1): 1–29.