

Bayesian Collaborative Bandits with Thompson Sampling for Improved Outreach in Maternal Health Program

Arpan Dasgupta
Google DeepMind
Bengaluru, India
arpandg@google.com

Gagan Jain
Google DeepMind
Bengaluru, India
jaingagan@google.com

Arun Suggala
Google DeepMind
Bengaluru, India
arunss@google.com

Karthikeyan Shanmugam
Google DeepMind
Bengaluru, India
karthikeyanvs@google.com

Milind Tambe
Harvard University, Google
DeepMind
Boston, USA
milindtambe@google.com

Aparna Taneja
Google DeepMind
Bengaluru, India
aparnataneja@google.com

ABSTRACT

Mobile health (mHealth) programs face a critical challenge in optimizing the timing of automated health information calls to beneficiaries. This challenge has been formulated as a collaborative multi-armed bandit problem, requiring online learning of a low-rank reward matrix. Existing solutions often rely on heuristic combinations of offline matrix completion and exploration strategies. In this work, we propose a principled Bayesian approach using Thompson Sampling for this collaborative bandit problem. Our method leverages prior information through efficient Gibbs sampling for posterior inference over the low-rank matrix factors, enabling faster convergence. We demonstrate significant improvements over state-of-the-art baselines on a real-world dataset from the world’s largest maternal mHealth program. Our approach achieves a 16% reduction in the number of calls compared to existing methods and a 47% reduction compared to the deployed random policy. This efficiency gain translates to a potential increase in program capacity by 0.5 – 1.4 million beneficiaries, granting them access to vital antenatal and post-natal care information. Furthermore, we observe a 7% and 29% improvement in beneficiary retention (an extremely hard metric to impact) compared to state-of-the-art and deployed baselines, respectively. Synthetic simulations further demonstrate the superiority of our approach, particularly in low-data regimes and in effectively utilizing prior information. We also provide a theoretical analysis of our algorithm in a special setting using Eluder dimension.

KEYWORDS

Thompson Sampling; Maternal Health; Multi-arm Bandits

ACM Reference Format:

Arpan Dasgupta, Gagan Jain, Arun Suggala, Karthikeyan Shanmugam, Milind Tambe, and Aparna Taneja. 2025. Bayesian Collaborative Bandits with Thompson Sampling for Improved Outreach in Maternal Health Program. In *Proc. of the 24th International Conference on Autonomous Agents*

and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 12 pages.

1 INTRODUCTION

Mobile health (mHealth) programs offer a powerful tool for delivering vital health information, but face a critical challenge: optimizing the timing of automated calls to maximize engagement. This is particularly important in maternal mHealth programs, which play a vital role in reducing maternal mortality rates - a key target within the WHO’s Sustainable Development Goals [3]. To ensure these programs achieve their full potential, automated calls must be strategically timed to achieve high pick-up rates, leading to improved health outcomes for mothers and infants.

This paper focuses on Kilkari [2], the world’s largest maternal mHealth program. Implemented nationwide across India by the Ministry of Health and Family Welfare in partnership with the NGO ARMMAN [1], Kilkari delivers critical maternal and child health information through automated voice calls throughout pregnancy and the post child birth period. Kilkari has served over 40 million mothers across India so far, with over 3 million active subscribers at any given time. The importance of listening to these voice messages has been shown to have significant impact on the health outcomes of mothers and babies [20], particularly among the most marginalised who have the most to benefit from this program, and have the least access to resources.

However, one major challenge faced by the program is that the pick-up rate of the calls is very low. This is largely due to the fact that different beneficiaries prefer to listen to these calls at different time slots and on different days due to practical constraints such as shared family phones, different working hours, household responsibilities, as well as network reliability, particularly in rural districts of India [15]. To address this, the program attempts sending the automated voice calls multiple times in a week, until a call is answered, with almost 50% of the economically weakest beneficiaries requiring more than 6 attempts on average [20] for a single message in a week and on average 23% beneficiaries being unreachable despite multiple attempts [16]. In fact, due to the scale of the program, consistent low listenership of the calls can even lead to beneficiaries being dropped from the program.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), A. El Fallah Seghrouchni, Y. Vorobeychik, S. Das, A. Nowe (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

⁰Code for experiments at <https://github.com/arpandg/ts-sgld>

Optimizing call timing is therefore essential to improve engagement, and maximize the dissemination of critical health information. Moreover, it could also help reduce critical bandwidth being spent heavily on retries, which would then enable scaling the outreach of the program to millions of more mothers across the country. While individual preferences would ideally inform call scheduling, Kilkari’s scale prohibits collecting individual time preferences, or demographic information that could help predict those preferences. This necessitates a robust, scalable solution to predict optimal call times based on limited information.

To this end, we formulate the call pick-up problem as a multi-user multi-armed bandit problem, where each time slot represents an arm, and pulling an arm corresponds to sending a call at that time. The reward is based on whether the call is answered. The bandit algorithm must learn which time slot to use for each user to maximize the probability of answering. Current state-of-the-art technique [25] formulated this problem as collaborative bandits with a low-rank assumption on user preferences. This low-rank assumption is justified by the observation that groups of users exhibit similar preferences in practice (Figure 1 shows an example of the low-rank pick-up problem). This work attempted to collaborate across users to quickly learn their preferences. To this end, the authors employed offline matrix factorization with Boltzmann exploration [11]. However, this approach is heuristic, lacks theoretical guarantees, and performs sub-optimal exploration, leading to worse regret in practice. Furthermore, it falls short when there are new users joining the program (as in the real world) as it requires a large number of samples to learn the preferences, hence delaying the inference of optimal slots which further increases the risk of drop off from the program. It also lacks the ability to incorporate any available prior information.

In this paper, we propose a Bayesian formulation for the collaborative bandit problem utilizing Thompson Sampling (TS). Bayesian solutions [9] allow the use of priors to quickly converge to a solution even with limited data, which is crucial in this context, as ineffective exploration could lead to delayed time slot inference and increased drop-off risk. Thompson Sampling [29] is also shown [22] to be a very effective method and can provide much tighter bounds than the previously used Boltzmann exploration.

While TS is efficient in terms of regret, its exact implementation is computationally expensive due to the cost of posterior sampling. To address this, we develop a computationally efficient heuristic based on Stochastic Gradient Langevin Dynamics (SGLD) for posterior sampling in TS. Empirically, we show that our algorithm outperforms existing techniques in both cluster and general low-rank settings. We demonstrate its superiority using a real-world dataset from the Kilkari program [7], showing how it can significantly increase call engagement and pick-up rates. While we focus on Kilkari, our methods are applicable to a wide range of mHealth programs.

The key contributions of the paper is a novel Thompson Sampling based algorithm for collaborative multi-armed bandit problem using alternating SGLD for efficient posterior sampling to make the algorithm scalable for potential real world deployment (Section 4). We also provide the first (to our knowledge) *Eluder dimension* characterization for a special clustered setting of the low rank bandit problem with infinite arms (Section 5) which is a subset of

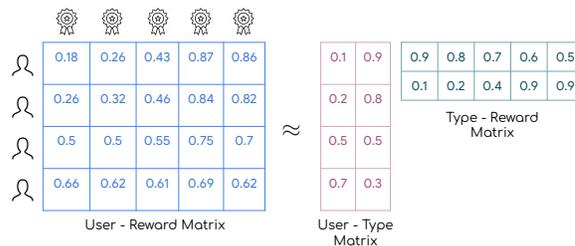


Figure 1: (Left) Example pick-up matrix—each entry such as 0.18 (top left) represents likelihood of a call being answered by the first beneficiary in the first time slot. (Right) Matrix decomposition into user X user type and user type X pickup rate probability matrices.

the general low-rank case. Eluder dimension is the principal complexity measure that characterizes Bayesian regret for Thompson sampling in general [26]. Finally, we demonstrate empirically the effectiveness of the approach on a time slot inference problem with a simulated and a real world dataset (Section 6) from the largest maternal mHealth program in the world.

The proposed algorithm exploits the collaborative nature of the problem to infer the optimal time slots to send automated voice messages to beneficiaries. We show that the proposed method reduces the number of call attempts needed to reach out to beneficiaries by a staggering **47 percent** compared to the current deployed system and **16 percent** in comparison to SOTA. Given that the program runs on a scale of millions of beneficiaries nationwide under limited budget constraints, a reduction of **16-47 percent** implies freeing up critical bandwidth **to potentially enrol 0.5-1.4 million more mothers**, disseminating critical health information and services to many more mothers from underprivileged communities who may otherwise have limited to no access to resources. We show a further 9 percent reduction in attempts compared to SOTA needed to reach future new enrolments in the program, leveraging the learnt matrix factorization, freeing up even more bandwidth on a regular basis.

Finally, due to the scale of the program and logistic constraints, sustained low listenership for several weeks in a row leads to beneficiaries being dropped off from the program. Despite this being an extremely hard metric to move, the proposed method is able to reduce these drop offs by 29 percent over a period of 4 months compared to the current deployment and by 7 percent compared to the SOTA. On the scale of a program operating for 3 million active subscribers currently, **an improvement of even 7 percent translates to 210,000 mothers being retained in the program** and continuing to benefit from access to valuable health information.

2 RELATED WORK

AI in Maternal Healthcare Limited resource allocation problems in maternal healthcare have previously been solved by restless multi-arm bandits [19, 21, 30]. The time slot selection problem using collaborative bandits was previously studied by [25].

Collaborative Multi-armed Bandits Multi-armed bandits are a highly studied and effective method for solving several resource allocation problems. Several methods such as phased elimination [17, 28], UCB [8], Thompson Sampling [4, 29] and Best-arm Identification [5, 13] have been studied in detail. The collaborative bandit

problem has recently garnered attention due to the widespread popularity of modern recommender systems [10, 12]. Under special conditions, several algorithms with strong bounds have been proposed [14, 24]. An algorithm which is applicable in approximate low rank was proposed by [25] which is close to our problem setting. In our work, we empirically compare against this work.

Offline Bayesian Matrix Factorization. While several methods have been proposed for Bayesian matrix factorization [22], the ones which utilize MCMC [27] are of particular interest to us. [6] proposes utilizing SGLD in a distributed manner using block partitioning to perform matrix factorization. While such methods have been applied in practice [18, 32], none of these methods extend the solution in a bandit setting that forms our focus.

Thompson Sampling & Eluder Dimension. Thompson Sampling (TS) is a widely used algorithm for bandit optimization, which maintains a prior distribution over unknown problem parameters and updates it sequentially as new data is observed [29]. It recommends actions by sampling from the posterior distribution. In a seminal work Agrawal and Goyal [4] showed that TS obtains optimal regret for Multi-Arm Bandits (MAB). Russo and Van Roy [26] further extended these results, deriving an upper bound for TS’s Bayesian-regret in terms of Eluder dimension, for general loss functions. Eluder dimension is a combinatorial quantity which measures the complexity of a function class. However, to the best of our knowledge, Eluder dimension has been characterized for very few function classes such as linear and quadratic [23].

3 PRELIMINARIES

Notation. We use \mathbb{R} to denote the set of real numbers. $\mathbb{R}^d, \mathbb{R}^{d_1 \times d_2}$ denote the sets of d -dimensional vectors, and $d_1 \times d_2$ dimensional matrices respectively. For a positive integer n , we use $[n]$ to denote the set $\{1, 2, \dots, n\}$. Δ_d denotes the probability simplex in \mathbb{R}^d .

Problem Formulation. Consider a mHealth program with N beneficiaries (users), M time slots (arms) and T rounds of calling. In the following, we interchangeably refer to the beneficiaries as users and the time slots as arms to describe our techniques in a more general fashion. Let $\Theta \in \mathbb{R}^{N \times M}$ be the (user, arm) reward matrix which is unknown to the system. Θ_{ij} denotes the expected reward of pulling arm j for user i . We assume that Θ can be decomposed as $\Theta = UV$, for some latent factors $U \in \mathbb{R}^{N \times C}, V \in \mathbb{R}^{C \times M}$, and $C \leq \min\{M, N\}$. Special cases of this model include the rank-1 setting (where $C = 1$), and the cluster setting (where each row of U is a one-hot encoded vector representing user cluster membership).

In each round $t \in [T]$, a user $u(t) \in [N]$ arrives at the system and is recommended an arm $a(t) \in [M]$. For ease of exposition, we assume the users arrive in a round-robin fashion; that is, $u(t) = t\%N$. At the end of round t , the system receives feedback in the form of reward from that user. Let $r_t \in [0, 1]$ be the observed noisy reward which satisfies: $\mathbb{E}[r_t] = \Theta_{u(t)a(t)}$. Note that $a(t)$ can depend on the history $H_t = \{u(s), a(s), r_s\}_{s=1}^{t-1}$. Finally, we let π be the prior distribution over the set of all rank C matrices (π encodes the prior knowledge about Θ). The goal of the learner is to optimize Bayesian

regret which is defined as

$$R(T; \pi) = \sum_{t=1}^T \mathbb{E} [\Theta_{u(t)*} - \Theta_{u(t)a(t)}], \quad (1)$$

where the expectation is taken over the prior distribution of Θ , internal randomness of the system, and Θ_{u*} denotes the reward of the best item for user u .

Thompson Sampling. As previously mentioned, TS is a widely used algorithm for bandit optimization. At any time step t , TS maintains a posterior distribution π_t over the set of all rank C matrices, representing its belief about the unknown reward matrix. TS samples a matrix $\hat{\Theta}^{(t)}$ from this distribution and recommends the arm with the highest predicted reward for the current user: $\arg \max_{j \in [M]} \hat{\Theta}_{u(t)j}^{(t)}$. Algorithm 1 describes this procedure.

Algorithm 1 Thompson Sampling

Input: number of rounds T , prior π
for $t = 1$ **to** T **do**
 Sample reward matrix: $\hat{\Theta}^{(t)} \sim \pi_t$
 Observe user $u(t)$, recommend $a(t) := \arg \max_{j \in [M]} \hat{\Theta}_{u(t)j}^{(t)}$
 Observe reward r_t
 Update posterior: $\pi_{t+1}(\Theta') \propto \mathbb{P}(\Theta' | H_{t+1})$
end for

In the next section, we develop a SGLD based algorithm for efficient posterior sampling in TS and demonstrate its superiority on a real-world maternal mHealth application (Section 6.3). In Section 5 we show that TS achieves optimal Bayesian regret guarantees for a cluster setting.

4 PROPOSED METHOD

4.1 Stochastic Gradient Langevin Dynamics (SGLD)

SGLD [31] is a Markov Chain Monte Carlo (MCMC) method which is used for Bayesian inference and is particularly beneficial in high-dimensional settings. SGLD is able to perform updates in batches, similar to SGD in optimization. The update equation is as follows

$$\Delta\Theta = \frac{\epsilon}{2} (\nabla \log p(\Theta) + \frac{N}{n} \sum_{i=1}^n \nabla \log p(X_i | \Theta)) + \eta \quad (2)$$

$$\eta \sim N(0, \epsilon)$$

where, Θ are the parameters that characterize the likelihood function of X (observed data). $p(\Theta)$ is the Bayesian prior over these parameters. After an initial set of updates, SGLD generates samples from the posterior distribution of the model parameters, allowing full Bayesian inference instead of just the point estimates.

In this section we go through the steps to apply Thompson Sampling to the low-rank collaborative bandit problem. The basic idea is to utilize SGLD [31] to be able to perform Bayesian Matrix Factorization as well as posterior sampling.

4.2 Bayesian Matrix Factorization

As described in Equation (2), SGLD allows us to compute the updates to the parameters in batches of data. In our case, we assume that the

matrices U, V are derived from parameters u, v . Also, in accordance with the assumption, $U_{i,c}$ is obtained as $\frac{e^{u_{i,c}}}{\sum_c (e^{u_{i,c'}})}$ and each element $V_{c,j}$ is filled with $\frac{1}{1+e^{v_{c,j}}}$. The parameters $u_{i,c}, v_{c,j}$ are sampled from the exponential distributions sampled from independent priors each of which is an exponential distribution with pre-decided parameters $\lambda_{i,c}$ and $\alpha_{c,j}$ respectively. Note that the complete set of parameters is $\Theta = \{u_{i,c}\} \cup \{v_{c,j}\}$.

In our setting the prior $p(\Theta)$ from Equation (2) is:

$$p(\Theta) = p(u, v) = \left(\prod_{i,c} \lambda_{i,c} e^{-\lambda_{i,c} u_{i,c}} \right) \left(\prod_{c,j} \alpha_{c,j} e^{-\alpha_{c,j} v_{c,j}} \right) \quad (3)$$

Coordinates of the gradient $\nabla_{\Theta} \log p(\Theta)$ are given by:

$$\nabla_{u_{i,c}} \log p(u, v) = \lambda_{i,c}, \quad \nabla_{v_{c,j}} \log p(u, v) = \alpha_{c,j} \quad (4)$$

Let us assume that the data X is composed of data points x where $x_d = x_{i,j}$ is the reward of the d th data point for user i for pulling arm j . The likelihood (second) term of Equation (2) is composed of the observed reward $x_{i,j}$ which is a mixture of Bernoulli random variables and is calculated as

$$Q = p(x_{i,j}|u, v) = \sum_c (p(c|u) p(x_{i,j}|c, v_j)) \quad (5)$$

$$p(c|u) = \frac{e^{u_{i,c}}}{\sum_k e^{u_{i,k}}}, \quad p(x_{i,j}|c, v_j) = \frac{x_{i,j} e^{v_{c,j}}}{(1 + e^{v_{c,j}})} + \frac{(1 - x_{i,j})}{(1 + e^{v_{c,j}})}$$

Here, $p(c|u)$ and $p(x_{i,j}|c, v_j)$ represent the probability of sampling archetype c given the user parameters and the Bernoulli likelihood of observing the Boolean reward $x_{i,j}$ given the rank c and the reward parameters respectively. Coordinates of matrix $\nabla_{\Theta} \log p(x_{i,j}|\Theta)$ are hence given by sum :

$$\nabla_{u_{i,c}} \log p(x_{i,j}|u, v) = \frac{1}{Q} \left(\frac{\delta U_{i,c}}{\delta u_{i,c}} \right) \cdot p(x_{i,j}|c, v_j) \quad (6)$$

$$\nabla_{v_{c,j}} \log p(x_{i,j}|u, v) = \frac{1}{P} \frac{e^{u_{i,c}}}{\sum_k e^{u_{i,k}}} \frac{(2x_{i,j} - 1)e^{v_{c,j}}}{(1 + e^{v_{c,j}})^2} \quad (7)$$

4.3 Thompson Sampling

Thompson Sampling [29], [22] is a popular solution for multi-armed bandits as it provides strong theoretical guarantees. TS requires maintaining a belief, acting according in the best possible way according to the belief and updating it as new evidence is collected. The most difficult step involved is the sampling step and it can be done effectively by MCMC methods like SGLD as they converge to the posterior distribution and generate samples from it. In our problem, we perform the updates to the parameters u, v and calculate $P = UV$ to be the final matrix, from which we choose the arm with the highest reward for each user.

4.4 Scaling the Algorithm

While SGLD is able to significantly speed up the posterior sampling process, the algorithm is still slow. In order to speed this up, we come up with a way to scale it in distributed settings. We observe that the updates to u are independent across users given v , ie. updates of one user u_i are independent from the updates of u_j assuming v to be fixed.

THEOREM 1. *The updates in parameters for one user are independent from the other users.*

$$P(u_k|v, u_{l \neq k}, X) = P(u_k|v, X) \quad (8)$$

PROOF. Given data X , parameters u, v , the likelihood function is:

$$L(X, u, v) = \prod_{k=1}^{|X|} \left((1 - X_{k(i,j)})(1 - u_i v_j) + (X_{k(i,j)} u_i v_j) \right) \quad (9)$$

where u_i is the i th row of u and v_j is the v th column of v .

We want to establish that the updates of the conditional posterior of the parameters of the k th user is independent of the parameters of the other users.

$$P(u_k|v, u_{l \neq k}, X) = \frac{L(u_k, r, \theta_{l \neq k}, X) p(u_k)}{\int L(u_k, v, u_{l \neq k}, X) p(u_i) du}$$

We can separate out the items in X where user k is involved.

$$L(u_i, v, u_{l \neq k}, X) = L(u, v, X) \quad (10)$$

$$= \prod_{l=1}^N \left((1 - X_{l(i,j)})(1 - u_i v_j) + (X_{l(i,j)} u_i v_j) \right) \quad (11)$$

$$= \left((1 - X_k)(1 - u_i v_j) + (X_k \theta_i v_j) \right) \quad (12)$$

$$\prod \left((1 - X_{l \neq k})(1 - u_i v_{j'}) + (X_{l \neq k} u_i v_{j'}) \right) \quad (13)$$

The second term in the product comes out in both numerator and denominator (as it is not dependent on u_k) and cancels out. Thus we are left with

$$P(u_k|v, u_{l \neq k}, X) = \frac{L(u_k, v, X_k) p(u_k)}{\int L(u_k, v, X_k) p(u_k) du} = P(u_k|v, X_k) \quad (14)$$

Where X_k is the data where k is involved. \square

Algorithm 2 and Algorithm 3 explain the full and alternating sampling SGLD methods. The complete algorithm is described by Algorithm 4.

Algorithm 2 SGLD with Full Sampling

Input: Batch Size n , Data X

Hyperparameters: Learning Rate ϵ

Parameters: λ, α

Initialize u and v

repeat

 Select batch x of size n from X .

 Calculate terms from Equation (6) and Equation (7)

 Update u, v using Equation (2).

until u, v converge

Return u, v

Theorem 1 allows us to compute updates in batches of users assuming v to be constant. The batch updates to u can be accumulated and then the updates to v can be computed. The updates to v also follow the conditional independence property, but the limiting factor in a majority of the cases is the number of users in u . These alternating updates are easier to perform than those in [6] as the recombination step is much easier.

Algorithm 3 SGLD with Alternating Sampling

Input: Batch Size n , Data X , User Blocks b
Hyperparameters: Learning Rate ϵ
Parameters: λ, α
Initialize u and v
repeat
 Select batch x of size n from X .
 for b_i **in** b **do**
 Calculate terms from Equation (6) for b_i
 Update u_{b_i} using Equation (2).
 end for
 Merge u_{b_i} to get u .
 Calculate terms from Equation (7)
 Update v using Equation (2).
until u, v converge
Return u, v

Algorithm 4 Proposed Algorithm TS-SGLD

Input: Time steps T , Samples per time step s .
Generate data X_0 from random users and time slots.
for $i = 1$ **to** T **do**
 if Sampling Method == *full* **then**
 u, v from Algorithm 2
 else if Sampling Method = *alternating* **then**
 u, v from Algorithm 3
 end if
 Calculate U, V
 Generate X_i with s samples using TS.
end for
Return U, V

5 THEORETICAL BOUNDS - CLUSTERED LOW RANK BANDITS

The general problem of low rank bandits in the Bayesian setup can be specified as follows. There is a space of low rank reward matrices that parameterized the problem given by: $\Omega_C \subset \mathbb{R}^{N \times D}$ such that $\text{rank}(R) = C, \forall R \in \Omega_C$. Here, we have N users and D is the dimension of the action space. Let the unknown reward matrix be R . We have a prior $\pi(R)$ over the space Ω .

Reward Likelihood Model: For user u and action vector $a \in \mathcal{A} \subset \mathbb{R}^D$, the reward model given R is given by:

$$Y(u, a) \sim R(u, \cdot)^T a + N \quad (15)$$

where N is i.i.d 1-sub-Gaussian random variable.

Bounded Parameter Space: We will assume that $\|R(u, \cdot)\|_2 < S, \forall R \in \Omega, \|a\|_2 < \gamma, \forall a \in \mathcal{A}$.

Bayesian Regret: Let user u_t arrive at time t and a_t be the action chosen according to some bandit algorithm given previous actions $(u_1, a_1), \dots, (u_{t-1}, a_{t-1})$ and the corresponding rewards $y_1 \dots y_{t-1}$. We are interested in Bayesian regret given by:

$$\mathbb{E}_{R \sim \pi, y(u_t, a_t) \sim Y(u_t, a_t) | R, u_t, \sim \text{Unif}[1:N]} [R(u_t, \cdot)^T a_t - R(u_t, \cdot)^T a^*(u_t)]$$

Here, $a^*(u_t) = \arg \max_{a \in \mathcal{A}} R(u_t, \cdot)^T a$

We analyze the Bayesian Regret for the natural Thompson Sampling algorithm with perfect Gibbs Sampling from the posterior at every step.

Thompson Sampling: At every time t , let the history of users, actions and rewards be denoted H_t given user u_t , compute posterior $\pi(R|H_t)$. Sample $\hat{R} \sim \pi(R|H_t)$. Play action $a_t = \arg \max_{a \in \mathcal{A}} \hat{R}(u_t, \cdot)^T a$.

The basis for our regret bounds is the seminal result of Russo and Van Roy [26] which upper bounds the Bayesian regret of TS in terms of Eluder dimension, a measure of the complexity of a model class. Specifically, for a model class $\Omega_C = \{UV : U \in \mathbb{R}^{N \times C}, V \in \mathbb{R}^{C \times D}\}$, the Eluder dimension is defined as the maximum length of a sequence of actions that consistently provide new information about the underlying function. Formally, an action u_t, a_t is said to be ϵ -dependent on previous action sequence $\{(u_1, a_1), \dots, (u_{t-1}, a_{t-1})\}$ if any pair of matrices $R, R' \in \Omega_C$ satisfying $\sqrt{\sum_{s=1}^{t-1} ((R_{u(s), \cdot} - R'_{u(s), \cdot})^T a_s)^2} \leq \epsilon$ also satisfies: $|(R'_{u(t), \cdot} - R_{u(t), \cdot})^T a_t| \leq \epsilon$. u_t, a_t is ϵ -independent of $\{(u_1, a_1), \dots, (u_{t-1}, a_{t-1})\}$ if it is not ϵ -dependent on it.

DEFINITION 1 (RUSSO AND VAN ROY [26]). *The ϵ -Eluder dimension $\text{dim}_E(\Omega_C, \epsilon)$ is the length of longest sequence of elements in $[C] \times [D]$ such that for some $\epsilon' \geq \epsilon$, every element is ϵ' -independent of its predecessors.*

Russo and Van Roy [26] show that the Bayesian regret of TS is upper bounded (upto logarithmic factors) by

$$\sqrt{\text{dim}_E(\Omega_C, T^{-2}) \log N(\Omega_C, T^{-2}, \|\cdot\|_\infty) T},$$

where $N(\Omega_C, \epsilon, \|\cdot\|_\infty)$ is the covering number of the model class. We see that ϵ -Eluder dimension is the key complexity measure that governs regret in bandit problems. Our main contribution is characterization of ϵ -Eluder dimension for the cluster model below which is a special case of low rank bandits with infinite arms. We describe the simpler model below.

Cluster Model with infinite arms: We further consider Ω_C with a clustering assumption that there are only C distinct rows in R and every user u is mapped to a cluster using a cluster assignment function $c : [N] \rightarrow [C]$. Therefore, one can reparameterize $\Omega_C = (c, R), c : [N] \rightarrow [C], R \in \mathbb{R}^{C \times D}$. $R(u, \cdot) = R(c(u), \cdot)$ for the reward model (15).

We will show that the ϵ -Eluder dimension for the above model is: $\text{dim}_E(\Omega_C, \epsilon) = O(2CD \log(1 + 2S/\epsilon^2) + CN)$. This is our key theoretical contribution.

While our main contribution is the ϵ -Eluder dimension characterization of the space of reward matrices for cluster model with infinite arms, we first establish ϵ -Eluder dimension of the finite arms case before proceeding to the general case as this is very insightful and simpler to understand.

Cluster Model with finite arms: We consider a simpler cluster problem where the action set $\mathcal{A} = \{e_1, e_2 \dots e_D\}$ consists of indicator vectors. We denote this as the cluster model with finite arms.

THEOREM 2. *In the cluster model, for C user-clusters with a total of N users and D finite arms D arms, the ϵ -Eluder dimension is at most $(2D + N)C$.*

PROOF SKETCH. We first take a specific realization of (c, R) and with respect to a sequence of actions look at the reward sequence which is ϵ Eluder. For every new action, some subset of alternate possible reward matrices get separated from this by taking a value ϵ far until only one is left. We argue that while separation happens you either learn something about the assignment function c for some $c(u)$ or learn something about entry in R . We then argue that corresponding Eluder subsequences where one or the other is learnt has lengths at most CN and CD respectively. Detailed formal argument is presented in the supplement. \square

THEOREM 3. *In the cluster model, for C user-clusters with a total of N users and D -dimensional infinite arm set \mathcal{A} , the ϵ -Eluder dimension is $O(2CD \log(1 + 2S/\epsilon^2) + CN)$.*

PROOF SKETCH. It follows the finite arms case except that we use a reduction to linear bandits to argue that the Eluder subsequence where new information about R is learnt cannot exceed $\tilde{O}(CD \log(1/\epsilon))$. Exact formal argument is given in the supplement. \square

THEOREM 4. *For the cluster model in the infinite arm case, we have:*

a) *the log-covering number $\log N(\Omega_C, \epsilon, \|\cdot\|_\infty) = O(CD \log(1/\epsilon) + N \log C)$ and,*

b) *$\dim_E(\Omega_C, \epsilon) = O(2CD \log(1 + 2S/\epsilon^2) + CN)$ (Theorem 3)*

Thus propositions 4 and 5 from [26] with $\epsilon = T^{-2}$ yields a Bayesian regret bound of:

$$\tilde{O}\left(CD\sqrt{T\left(1 + \frac{N}{D}\right)\left(1 + \log(C)\frac{N}{CD}\right)}\log(T)\right)$$

PROOF. Theorem 3 shows that $\dim_E(\Omega_C, \epsilon) = O(2CD \log(1 + 2S/\epsilon^2) + CN)$. In the clustered bandits case with infinite arms, the reward matrix has two components, $R \in \mathbb{R}^{C \times D}$ and an assignment function $c : [N] \rightarrow C$. If we want to cover space of matrix R in infinity norm, then we would exactly cover all assignment functions and there are C^N in number and for every R we would find an approximation which is ϵ far in the infinity norm. The latter ϵ cover would involve $O(\frac{1}{\epsilon})^{CD}$ matrices assuming a bounded space of $C \times D$ matrices. Therefore, the covering number is $O(\frac{1}{\epsilon})^{CD} * C^N$. This yields the log-covering number bound. Combining both by results in [26] gives the regret bound. \square

Remark: Assuming every user arrival (random uniform arrival) is treated as an independent linear bandit, regret would be roughly $\tilde{O}(D\sqrt{NT})$. When $N \approx D$ (i.e. N and D are comparable), our regret bound is effectively $\tilde{O}(CD\sqrt{T})$ ignoring log and constant factors where only the number of clusters and D matters. Usually C is a constant while N and D are very high dimensional. In this case, our regret bound gives a polynomial advantage over treating each user as an independent bandit. When $N \gg D$, we point out that we have an extra factor $\sqrt{(1 + \log C \frac{N}{CD})}$ in regret that appears primarily because of the log covering number dependence in regret. We believe that a tighter analysis connecting regret and Eluder dimension is needed in this specific clustering setting to resolve this.

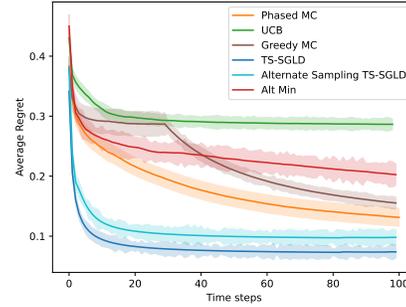


Figure 2: Regret for the low-rank case on simulated data: Average regret for different methods averaged over 15 random matrices. Results are on a 1000 users and 20 arms matrix with 4 user types. Every time step adds 1000 samples.

6 EXPERIMENTS

6.1 Baselines

We show the efficacy of the proposed algorithm **TS-SGLD** (Algorithm 4) by comparing against other methods performing low-rank collaborative bandits. We compare our method against the SOTA Phased Matrix Completion (Phased MC) and Greedy Matrix Completion (Greedy MC) from [25]. Other baselines include Alternating Linear Bandits (AltMin) from [12] and LATTICE from [24], where the latter is specifically for comparison on the simulated cluster setting. We also consider UCB on individual arms and a random baseline which is currently deployed by the NGO for the real world maternal health application. Results are shown for the best fine-tuned hyperparameters for all baselines to give a fair comparison.

6.2 Simulated Data

The simulated dataset is created by generating random matrices of size $N \times M$ with rank C by multiplying two random matrices U and V of shape $N \times C$ and $C \times M$ respectively and adding Gaussian noise to it. This is followed by a normalization step which scales the values to probabilities between $[0, 1]$.

$$P = \text{normalize}(U \cdot V + N(\mu, \sigma^2))$$

The matrices U, V are randomly generated with each element being from the range $[0, 1]$ and μ and σ are chosen to be 0.5 and 0.1 respectively. The *normalize* operation maps the smallest and highest values to 0 and 1 and scales the other values accordingly. C is set to 4 for these experiments. Uniform prior of 0.5 is applied to all parameters in Equation 3.

Figure 2 compares regrets of different methods for 1000 users and 20 arms over 15 randomly generated low-rank matrices. UCB gives the highest regret since it cannot leverage the collaborative structure of the arm rewards. Due to the low-data setting, **Greedy MC** takes more exploration steps to catch up with **PhasedMC**. **GreedyMC** is primarily an explore-then-exploit approach which is prohibitive since sub-optimal choices at the start can be costly, especially in the subsequent real-world application. **TS-SGLD** with full sampling performs about **65%** better than **PhasedMC**. **TS-SGLD** with alternate sampling has higher regret than **TS-SGLD**

with full sampling since alternating sampling leads to a noisier convergence to the posterior, but still has lower regret compared to all other baselines. Similar results are also shown for the special cluster setting described in Section 5 in Figure 5a in the Appendix. LATTICE [24] is known to provide optimal bounds for the cluster case in theory. However in practicality, in the low data regimes where there is a small amount of data coming in at each time step, the regret of LATTICE is worse than **TS-SGLD** with full sampling by 15 percent in final cumulative regret.

6.3 Real World Dataset from Maternal mHealth program

Dataset: The problem of optimal time slot selection for beneficiaries in a mHealth program can be formulated as a collaborative bandits problem [25]. The dataset is created from historic data from the Kilhari program where anonymized call logs provide information about which user was called at which time slot, and whether the call was answered to provide a groundtruth pickup matrix (consisting of **200,000** users observed over one year for 7 time slots daily between *8am* and *8pm*) and a test dataset corresponding to a different set of **1000** users. For the experiments, timeslots are discretized into 7 and 14 slots where 14 slots correspond to different slots for weekdays and weekends. 14 slots increases the problem complexity, however it is more representative of this domain given that many users have limited access to shared phones on weekdays and may have different preferences for weekends.

Prior: Prior information is obtained from previous iterations of the calling program. It was observed that the real world dataset for a significant number of users contains zero entries in the pickup matrix i.e. zero pick-up rates across most time slots. This was provided as a prior on the V matrix on the second term in Equation (3). No priors were provided on the U matrix, since there is no demographic or any other meaningful information available about the users to provide meaningful priors on the user type categorization matrix. However, the method allows for priors to be incorporated on both U and V matrices if such information were available.

The method requires that we choose an appropriate rank C which fits the data reasonably. In all of our experiments for this dataset, $C = 5$ was chosen, based on experiments shown in Figure 6 in the Supplementary material. Users are split into different buckets based on largest pick-up probability across time slots to analyze the performance of the algorithms finely. Beneficiaries with highest pick-up rate (across any time slot) below 0.2 and above 0.8 are classified into the low and high pick-up rate groups respectively. The remaining users are classified into the middle pick-up rate group. On a dataset of 1000 users, 488 users fall in this category.

Impact on Regret: Figure 7b (in supplementary) compares average regret on the medium listenership bucket of users across different baselines, and Figure 7a (in supplementary) shows regret across all users. **TS-SGLD** performs particularly well in low-data settings such as these. **TS-SGLD** with prior, is able to effectively utilize prior information and outperforms **TS-SGLD** without prior by 8 percent as well as **PhasedMC** and **AltMin** by 14 percent and 21 percent respectively. Despite the extremely low pickup and high pickup users being harder to impact, **TS-SGLD** improves these by 4 and 8 percent respectively compared to SOTA.

Impact on re-attempts: Since a lot of bandwidth is currently spent on re-attempts to increase likelihood of pickup across users, we also measure the expected number of attempts required by the different methods. The program is known to implement a constraint of ≤ 9 calls [20]. The average number of attempts made for connecting a call is 5 [20]. Figure 3a shows the average number of attempts made for the beneficiaries in the middle pick-up rate group (49% of beneficiaries) reduces by 16% compared to the SOTA and 46% compared to the current deployed system.

Weekday vs Weekend slots: Figure 3c shows the number of attempts reduces significantly compared to the SOTA and non-collaborative baselines, when 14 slots are used instead of 7 slots to take into account additional preference for weekday vs weekends.

Impact on new enrollments: Since the program enrolls new beneficiaries at regular intervals, Figure 7e and Figure 3d show the performance of the algorithm for new users after the algorithm has been running on an existing cohort for 10 timesteps for both the algorithms. The regret achieved is significantly lower for **TS-SGLD** compared to **PhasedMC** by 18 percent and translates to a further 9% reduction in the number of attempts to connect calls to the new users. In fact, due to the UV factorization, the algorithm only has to learn the U matrix entries for the new users, and can leverage the already known V matrix to infer slots sooner. Since these mHealth programs suffer from a high rate of drop off [7] due to prolonged low listenership, it is important not only to learn optimal time slots accurately but also quickly with a few number of samples to avoid risk of drop off from the program.

Impact on Drop offs: Currently, beneficiaries with low listenership ($< 25\%$ of message length) for 6 consecutive weeks or low listenership for 9 weeks within a 16 week period are dropped off from the program. It is therefore imperative to send calls at a convenient time to both boost listenership and retain beneficiaries in the program to ensure continued access to critical health information. Despite this being a very strict metric and considerably harder to impact, **TS-SGLD** manages to reduce dropoffs by 7 percent compared to the SOTA and by 29 percent compared to the current deployed system (Figure 4).

Combined Pickup and Engagement data: Similar to [25], pick-up and engagement data was combined to minimize the regret for engagement. A call is said to be *engaged* if listenership is greater than 25% of the message length. **TS-SGLD** on concatenated pickup and engagement data shows 17 percent improvement (Figure 3b) in re-attempts over SOTA on medium-bucket listenership users.

Simulated Data vs Real World Data Gap: The difference in performance of the algorithm between the simulated and real-world data is due to the increased amount of unstructured noise in real datasets. Furthermore, the low-rank assumption in the real-world dataset does not hold completely, and in fact the matrix is approximately low-rank, as discussed in Supplementary Section A.

7 CONCLUSION AND FUTURE WORK

We presented an algorithm to solve a collaborative bandits problems for low rank matrices using Thompson sampling. The algorithm was run on time slot inference problem inspired by the real world dataset obtained from the largest maternal mHealth program in the world. The proposed method showed a significant reduction in the number

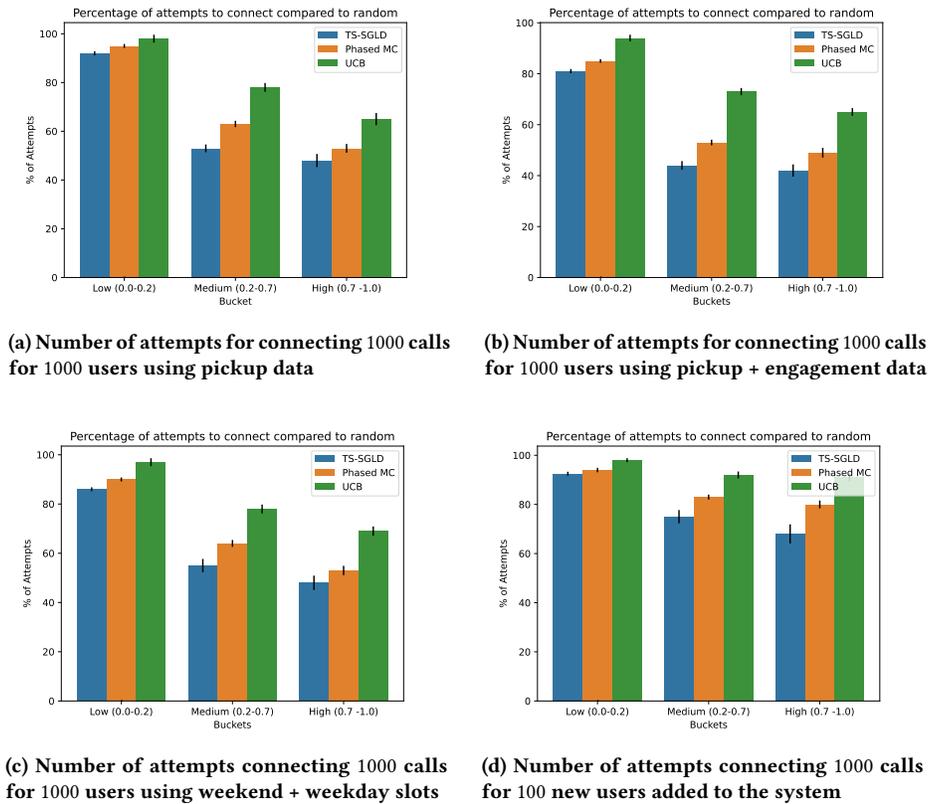


Figure 3: Number of attempts needed to reach out to beneficiaries in the real-world ARMMAN dataset across 3 listenership buckets. All plots are relative to the deployed random baseline capped at 100%.

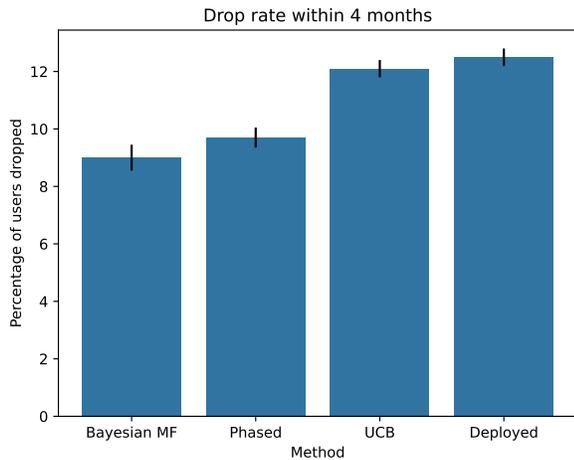


Figure 4: Percentage of dropoffs over a 4 month period. Dropoffs happen when the engagement goes below 25% for 6 weeks consecutively or 9 weeks in a 12 week period.

of call attempts needed to reach out to beneficiaries by 47 percent compared to the current deployment and 16 percent compared to the SOTA matrix completion methods, and a further 9 percent for new enrolments in the program leveraging the already inferred user type to slot preference information from matrix factorization. Reduced attempts free up critical bandwidth, to enable the program operating under budget restrictions to potentially enroll 0.5 to 1 million more mothers (assuming on average 5 attempts per voice message) into the program. The method also led to a reduction in drop offs by 7 and 29 percent compared to the SOTA and the current deployment, which can effectively enable retaining 0.2 to 1 million mothers into the program, ensuring continued access to critical health information for marginalised communities with limited access to resources. We further strengthened our approach by utilizing both the pickup and engagement data. To the best of our knowledge, we are also the first to provide Eluder dimension based analysis for Thompson Sampling for the cluster case which is a subset of the general low rank case. While the proposed method can incorporate priors on both user type and user preference matrices, we currently incorporate priors learnt from previous calling data only on the latter due to lack of demographic information about beneficiaries.

REFERENCES

- [1] [n. d.]. Armman Home - ARMMAN - Helping Mothers and Children – armman.org. <https://armman.org>. [Accessed 31-05-2024].
- [2] [n. d.]. Kilkari - ARMMAN - Helping Mothers and Children – armman.org. <https://armman.org/kilkari/>. [Accessed 31-05-2024].
- [3] [n. d.]. SDG Target 3.1 Maternal mortality – who.int. <https://www.who.int/data/gho/data/themes/topics/sdg-target-3-1-maternal-mortality>. [Accessed 31-05-2024].
- [4] Shipra Agrawal and Navin Goyal. 2012. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*. JMLR Workshop and Conference Proceedings, 39–1.
- [5] Shubhada Agrawal, Sandeep Juneja, and Peter Glynn. 2020. Optimal δ -Correct Best-Arm Selection for Heavy-Tailed Distributions. In *Algorithmic Learning Theory*. PMLR, 61–110.
- [6] Sungjin Ahn, Anoop Korattikara, Nathan Liu, Suju Rajan, and Max Welling. 2015. Large-scale distributed Bayesian matrix factorization using stochastic gradient MCMC. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 9–18.
- [7] ARMMAN. 2023. Kilkari. <https://armman.org/kilkari/>
- [8] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47 (2002), 235–256.
- [9] Jasmin Praful Bharadiya. 2023. A Review of Bayesian Machine Learning Principles, Methods, and Applications. *International Journal of Innovative Science and Research Technology* 8, 5 (2023), 2033–2038.
- [10] Guy Bresler, Devavrat Shah, and Luis Filipe Voloch. 2016. Collaborative filtering with low regret. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*. 207–220.
- [11] Nicolò Cesa-Bianchi, Claudio Gentile, Gábor Lugosi, and Gergely Neu. 2017. Boltzmann exploration done right. *Advances in neural information processing systems* 30 (2017).
- [12] Hamid Dadkhahi and Sahand Negahban. 2018. Alternating linear bandits for online matrix-factorization recommendation. *arXiv preprint arXiv:1810.09401* (2018).
- [13] Aurélien Garivier and Emilie Kaufmann. 2016. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*. PMLR, 998–1027.
- [14] Prateek Jain and Soumyabrata Pal. 2022. Online low rank matrix completion. *arXiv preprint arXiv:2209.03997* (2022).
- [15] Bashingwa JJH, Mohan D, Chamberlain S, Arora S, Mendiratta J, Rahul S, Chauhan V, Scott K, Shah N, Ummer O, Ved R, Mulder N, and LeFevre AE. 2021. Assessing exposure to Kilkari: a big data analysis of a large maternal mobile messaging service across 13 states in India. *BMJ Glob Health* (2021).
- [16] Arshika Lalan, Shresth Verma, Kumar Madhu Sudan, Amrita Mahale, Aparna Hegde, Milind Tambe, and Aparna Taneja. 2023. Analyzing and Predicting Low-Listenership Trends in a Large-Scale Mobile Health Program: A Preliminary Investigation. *arXiv preprint arXiv:2311.07139* (2023).
- [17] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit algorithms*. Cambridge University Press.
- [18] Gai Li, Qiang Chen, et al. 2016. Exploiting explicit and implicit feedback for personalized ranking. *Mathematical Problems in Engineering* 2016 (2016).
- [19] Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. 2022. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 12017–12025.
- [20] Diwakar Mohan, Kerry Scott, Neha Shah, Jean Juste Harrison Bashingwa, Arpita Chakraborty, Osama Ummer, Anna Godfrey, Priyanka Dutt, Sara Chamberlain, and Amnesty Elizabeth LeFevre. 2021. Can health information through mobile phones close the divide in health behaviours among the marginalised? An equity analysis of Kilkari in Madhya Pradesh, India. *BMJ Global Health* 6, Suppl 5 (2021), e005512.
- [21] Vineet Nair, Kritika Prakash, Michael Wilbur, Aparna Taneja, Corinne Namblard, Oyindamola Adeyemo, Abhishek Dubey, Abiodun Adereni, Milind Tambe, and Ayan Mukhopadhyay. 2022. Adviser: Ai-driven vaccination intervention optimiser for increasing vaccine uptake in nigeria. *arXiv preprint arXiv:2204.13663* (2022).
- [22] Shinichi Nakajima and Masashi Sugiyama. 2011. Theoretical analysis of Bayesian matrix factorization. *The Journal of Machine Learning Research* 12 (2011), 2583–2648.
- [23] Ian Osband and Benjamin Van Roy. 2014. Model-based reinforcement learning and the eluder dimension. *Advances in Neural Information Processing Systems* 27 (2014).
- [24] Soumyabrata Pal, Arun Sai Suggala, Karthikeyan Shanmugam, and Prateek Jain. 2023. Optimal algorithms for latent bandits with cluster structure. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 7540–7577.
- [25] Soumyabrata Pal, Milind Tambe, Arun Suggala, and Aparna Taneja. 2024. Improving Mobile Maternal and Child Health Care Programs: Collaborative Bandits for Time slot selection. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 1540–1548.
- [26] Daniel Russo and Benjamin Van Roy. 2013. Eluder dimension and the sample complexity of optimistic exploration. *Advances in Neural Information Processing Systems* 26 (2013).
- [27] Ruslan Salakhutdinov and Andriy Mnih. 2008. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the 25th international conference on Machine learning*. 880–887.
- [28] Aleksandrs Slivkins et al. 2019. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning* 12, 1-2 (2019), 1–286.
- [29] William R Thompson. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3-4 (1933), 285–294.
- [30] Shresth Verma, Gargi Singh, Aditya S Mate, Paritosh Verma, Sruthi Gorantala, Neha Madhiwalla, Aparna Hegde, Divy Hasmukhbhai Thakkar, Manish Jain, Milind Shashikant Tambe, et al. 2023. Deployed SAHEL: Field Optimization of Intelligent RMAB for Maternal and Child Care. In *Innovative Applications of Artificial Intelligence (IAAI)*.
- [31] Max Welling and Yee W Teh. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. Citeseer, 681–688.
- [32] Chihao Zhang, Yang Yang, Wei Zhou, and Shihua Zhang. 2020. Distributed bayesian matrix decomposition for big data mining and clustering. *IEEE Transactions on Knowledge and Data Engineering* 34, 8 (2020), 3701–3713.