SAHELI for Mobile Health Programs in Maternal and Child Care: Further Analysis

Shresth Verma[1*], **Gargi Singh**[1*], **Aditya Mate**[1 2†], **Neha Madhiwalla**[3], **Aparna Hegde**[3], **Divy Thakkar**[1], **Manish Jain**[1], **Milind Tambe**[1], **Aparna Taneja**[1]

[1]Google Research India [2]Harvard University [3]ARMMAN

{vermashresth, gargisingh}@google.com, aditya_mate@g.harvard.edu, {neha, aparnahegde}@armman.org, {dthakkar, manishjn, milindtambe, aparnataneja}@google.com

Abstract

Underserved communities face critical health challenges due to lack of access to timely and reliable information. Non-governmental organizations are leveraging the widespread use of cellphones to combat these healthcare challenges and spread preventative awareness. The health workers at these organizations reach out individually to beneficiaries; however such programs still suffer from declining engagement.

We have deployed SAHELI, a system to efficiently utilize the limited availability of health workers for improving maternal and child health in India. SAHELI uses the Restless Multi-armed Bandit (RMAB) framework to identify beneficiaries for outreach. It is the *first deployed application* for RMABs in public health, and is already *in continuous use* by our partner NGO, ARMMAN. We have already reached $\sim 100K$ beneficiaries with SAHELI, and are on track to serve 1 million beneficiaries by the end of 2023. This scale and impact has been achieved through multiple innovations in the RMAB model and its development, in preparation of real world data, and in deployment practices; and through careful consideration of responsible AI practices. Specifically, in this paper, we describe our approach to learn from past data to improve the performance of SAHELI's RMAB model, the real-world challenges faced during deployment and adoption of SAHELI, and the end-to-end pipeline. Additionally, we showcase the characteristics of beneficiaries who benefit the most from SAHELI.

Introduction

Mobile health (mHealth) programs, that leverage the widespread use of cellphones, are a crucial resource for bridging information inequities for underserved and marginalized communities in the global south (Tshikomana and Ramukumba 2022; Gupta et al. 2022), especially in areas such as public health and social services where access to authoritative information is unevenly distributed. Many non-governmental organizations (NGOs) periodically send automated voice messages to improve health outcomes of beneficiaries. However, in spite of high adoption, adherence is a key challenge in public health information programs (ARMMAN 2019; Jakob et al. 2022; Eysenbach 2005; Meyerowitz-Katz et al. 2020). NGOs often employ live service calls made by health workers to boost engagement via encouragement or through logistic changes requested by beneficiaries. However, given the comparatively large number of potential beneficiaries, it is important to maximally utilize the limited availability of health workers, and thus it is crucial to identify the best recipients for such service calls.

States Covered in India	19
Partner NGOs	40
Partner Hospitals	97
Health Workers Trained	235K
Beneficiaries	27.2M

Scale of ARMMAN

Figure 1: A beneficiary receiving preventive health information

While AI models can help health workers in optimizing their service calls, deploying these models in the context of mHealth programs for underserved communities presents unique challenges. First, available data is sparse and skewed (because data is necessarily limited from small numbers of service calls). Second, NGOs are constrained by a very limited compute budget. Third, responsible deployment of the AI models is particularly important in such settings.

In this paper, we show how we address these research challenges in our deployed AI model – a Restless Multi-Armed Bandits (RMAB) model – together with our NGO partner ARMMAN (ARMMAN 2008) to improve the quality of service of their mHealth program focusing on maternal and child care in India [1]. India suffers from high maternal and neonatal mortality rates (Meh et al. 2022; World Health Organization (WHO) 2020), and ARMMAN (ARMMAN 2008) runs one of the largest mHealth programs in this domain in India. Our system, SAHELI (System for Allocating Healthcare-resources Efficiently given Limited Interventions), is the result of deep partnership of an inter-

[*]Joint first author

[†]Work done during an internship at Google Research India

[1]This paper builds on and provides additional analysis over the real-world results of Verma et al. (2023)

disciplinary team of researchers. SAHELI (meaning 'female friend' in Hindi) is designed to assist, rather than substitute, health workers in their normal workflow. The key contributions of deployed SAHELI are:

- SAHELI includes the first deployed application of RMABs for public health, and it is continuously in use by our partner NGO ARMMAN.
- A key novelty of the deployment is that it both predicts RMAB model parameters and computes optimal policies; in contrast with most past research that has focused on computing optimal policies. To that end, we provide an improved and robust machine learning prediction framework by performing model selection and evaluation of real-world RMAB systems.
- We deployed SAHELI on cloud infrastructure with an emphasis on frugality throughout the end-to-end pipeline given the resource constraints of the NGO partner.
- We present an impact analysis of SAHELI showcasing the characteristics of beneficiries who benefit the most from our system.

SAHELI has been developed as a platform, with the ability to be scaled to more NGOs in more domains. Our source code and data dictionary are available on Github[2]. For details on the ethics and data use, please refer to the appendix.

Related Work

While several works in the healthcare domain have studied patient adherence for diseases like HIV (Tuldrà et al. 1999), cardiac problems (Corotto et al. 2013) and tuberculosis (Killian et al. 2019; Pilote et al. 1996), these largely focus on building machine learning classifiers to predict future adherence to prescribed medication. With such models, the pool of beneficiaries flagged as 'high-risk' can itself be very large. Furthermore, the one-shot predictions of these models fail to capture the sequential decision making aspect of the problem. Other approaches that consider sequential decision making challenges, such as Pollack et al. (2002); Liao et al. (2020) adopt reinforcement learning techniques to build personalized health monitors that can send timely notifications or activity suggestions to users. However, these models assume notifications can be sent at will, and as such, do not address the challenge of limited service call resources.

Alternatively, RMABs have seen significant theoretical investigation, motivated by resource allocation challenges, such as in anti-poaching patrols (Qian et al. 2016), multi-channel communication (Liu and Zhao 2010), sensor monitoring and machine maintenance tasks (Glazebrook, Ruiz-Hernandex, and Kirkbride 2006). While they provide important contributions, none of these works have seen a real world deployment, and most have not been field tested.

Key reasons for the lack of RMAB deployment are their significant computational and data requirements. For example, just the optimization problem of computing the optimal allocation π, while assuming the transition parameters $\mathcal{P}$ are available, is already known to be PSPACE-hard (Papadimitriou and Tsitsiklis 1999). Furthermore, in the real world, these transition parameters are not just unknown but also hard to infer for real beneficiaries enrolling with ARMMAN and other similar health programs, as they come with no historical transition data. Despite such difficulties, our work is the first to deploy RMABs in tackling a real-world maternal healthcare task via frugal design choices discussed below.

[2]https://github.com/armman-projects/SAHELI

Problem Introduction

ARMMAN is a non-governmental nonprofit organization based in India, focused on improving maternal and child health outcomes among underserved and underprivileged communities (ARMMAN 2008). Their flagship program, 'mMitra', is a mHealth service that aims to leverage the extensive cellphone penetration in India to send out critical preventive health information to expectant or new mothers via automated voice messages. A large fraction ($\sim 90\%$) of mothers in the mMitra program are below the World Bank international poverty line (World Bank 2020). Despite the acute economic disadvantages faced by these mothers, such automated voice messages prove to be a feasible mode of information dissemination at scale, thanks to the wide accessibility of low-cost phones.

After enrollment into the mMitra mHealth program, beneficiaries receive 1-2 minute voice messages with health information according to the beneficiary's gestational age or age of the infant. Unfortunately, despite the proven effectiveness of this information program in improving maternal health outcomes, ARMMAN often sees dwindling engagement rates among beneficiaries, including frequent dropouts. Around 22% of beneficiaries drop out of the program after just 3 months. To counter this issue, ARMMAN leverages health workers that place live service calls (phone calls) to a limited number of beneficiaries on a weekly basis to encourage beneficiaries' participation, address requests/ complaints, and attempt to prevent engagement drops. This raises the key question of deciding which beneficiaries to pick for live service calls in order to improve engagement rates among the beneficiaries.

Restless Multi-Armed Bandits (RMAB)

The Restless Multi-Armed Bandits (RMABs) model was first introduced by Whittle (1988) to address limited resource allocation problems, but has not received much attention in terms of real-world deployments. An RMAB consists of a set of N arms, where each arm is associated with a two-action MDP (Puterman 2014). An MDP $\{\mathcal{S}, \mathcal{A}, r, P\}$ consists of a set of states $\mathcal{S}$, a set of actions $\mathcal{A}$, a reward function $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$, and a transition function P, where $P^{\alpha}_{s,s'}$ is the probability of transitioning from state s to s' when action α is chosen. The reward function in our setup is given as $r(s, \alpha, s') = s'$. An MDP policy $\pi : \mathcal{S} \mapsto \mathcal{A}$ maps to the choice of action to take at each state. The long-term discounted reward for a policy π, starting from state $s_0 = s$ is defined as $R^{\pi}_{\gamma}(s) = E\left[\sum_{t=0}^{\infty} \gamma^t r(s_{t+1}) | s_0 = s\right]$ where $s_{t+1} \sim P^{\pi(s_t)}_{s_t, s_{t+1}}$ and $\gamma \in [0, 1)$ is the discount factor. The total reward in the RMAB is defined as the sum of the total rewards accrued by individual arms of the RMAB.

In the setup we consider, each arm of the RMAB models a beneficiary enrolled with ARMMAN, who can be in one of two states $\mathcal{S} = \{0, 1\}$ (corresponding to 'Not Engaging (NE)' and 'Engaging (E)' respectively). Engagement in our setup was defined in consultation with the subject matter experts at ARMMAN: we define a beneficiary as engaged when she listens to at least one call in a week for more than 30 seconds. The action space for each arm consists of two actions, $\mathcal{A} = \{0, 1\}$, where $1(0)$, typically called the active (passive) action, refers to selecting (not selecting) the beneficiary for the live service call. Beneficiaries may transition from say their E state to NE state (or other transitions) from one week to the next week based on their transition probabilities defined on passive or active actions. The planner's goal is to select actions on arms (deliver live service calls) so as to maximize the total reward, i.e. number of beneficiaries in the engaged state, accrued by the RMAB. However, the budget constraint demands that the planner can choose no more than k arms ($k \ll N$) for the active action at any given timestep, i.e., no more than k live service calls per week.

The dominant technique for solving RMABs uses the Whittle Index heuristic (Whittle 1988), which is shown to have asymptotic optimality under some conditions (Weber and Weiss 1990), and to provide excellent performance in practice (Qian et al. 2016). Whittle indexes are formulated using the idea of passive subsidy, and informally rank arms so as to choose the top k, based on how attractive it is for a planner to activate each arm. For computing Whittle index, we use binary search algorithm from Qian et al. (2016)

Previous Study: Our previous study conducted in April 2021 (Mate et al. 2022) is the first to present real-world service quality improvement using RMABs in the context of mMitra program. This study tested an RMAB-based policy against two baselines of interest, and showed RMAB outperforming its competitors. The study spanned 7 weeks and included $23,003$ real-world beneficiaries who were distributed in three groups corresponding to the RMAB policy, Round Robin (RR) and Current Standard of Care (CSOC). Whereas RR corresponds to a non-AI heuristic for systematically calling beneficiaries, CSOC did not call any individuals. The results from this pilot study are shown in Table 1.

Improvements	**RMAB over CSOC**	**RMAB over RR**	**RR over CSOC**
% reduction in total beneficiary engagement drops	32.0%	28.3%	5.2%
p-value	0.044	0.098	0.740

Table 1: RMABs demonstrate statistically significant superior performance when compared against other non-AI approaches, namely Current Standard of Care (CSOC) and Round Robin (RR), as shown by Mate et al. (2022).

The pilot results demonstrated that the RMAB method cuts $\sim 30\%$ of the beneficiary engagement drops experienced by the other groups. Furthermore, whereas RMAB achieves statistically significant improvement against CSOC ($p < 0.05$) and RR ($p < 0.1$), RR fails to achieve any statistically significant improvement over CSOC. This key result forms the basis of relying on RMAB-based strategy over other non-AI strategies as a basis of SAHELI. In this paper, we describe the journey from this initial study to the final deployment. Whereas we use the same overall RMAB learning and optimization approach, we made multiple changes to provide significant enhancements that reduce data anomalies and improve computational performance of this RMAB-based strategy. Additionally, our deployed cloud application now automates the data exchange process with the NGO's systems while requiring minimal compute resources to be feasibly handled by the NGO. We now describe the end-to-end SAHELI system.

Deploying SAHELI

We now introduce SAHELI and its architecture. We begin by discussing the different components, and follow that up with the description of the AI pipeline. We then discuss the frugal design choices – both in modeling and infrastructure – that were required to finalize the deployment.

System Architecture

We first describe all the interactions within SAHELI's ecosystem (refer Figure 2). The health workers in the field periodically register beneficiaries through door-to-door visits or at the hospitals (step 1). The socio-demographic data such as age, language, income range, as well as the information on gestational age is then entered into the database maintained by ARMMAN (step 3). Automated voice messages tailored to the beneficiaries' gestational age are sent with the help of a telecommunication provider (step 4). The meta-data of the outcome such as duration of the call, failure reason etc, is also pushed to ARMMAN's database . As beneficiaries' engagement with the voice messages diminishes over time, live service calls are made by ARMMAN to encourage beneficiaries to engage with the program (step 10). However due to limited resources on the NGO's side, only a limited number of live service calls can be made each week. The AI pipeline predicts which beneficiaries would benefit most from receiving a service call in any given week. This list of beneficiaries is then generated at the start of each week and distributed across health workers in an automated fashion as shown on Figure 2 in steps through 2-9.

Step 8 in Figure 2 shows the generation of the list of beneficiaries that should be intervened in the given week using the AI pipeline. This list is ingested in ARMMAN's cloud databases, which serve as the back-end of a client mobile application (screenshot provided in Figure 2) used by the health workers. This client application randomly distributes the list of scheduled service calls among health workers based on their weekly availability. An illustrative screenshot (not real beneficiary) is also shown in Figure 2. The calls are made through the week (step 10) with a maximum of 3 call attempts to the same beneficiary. All the beneficiaries in the generated list receive the aforesaid service calls. The model is currently providing services to beneficiaries enrolling at

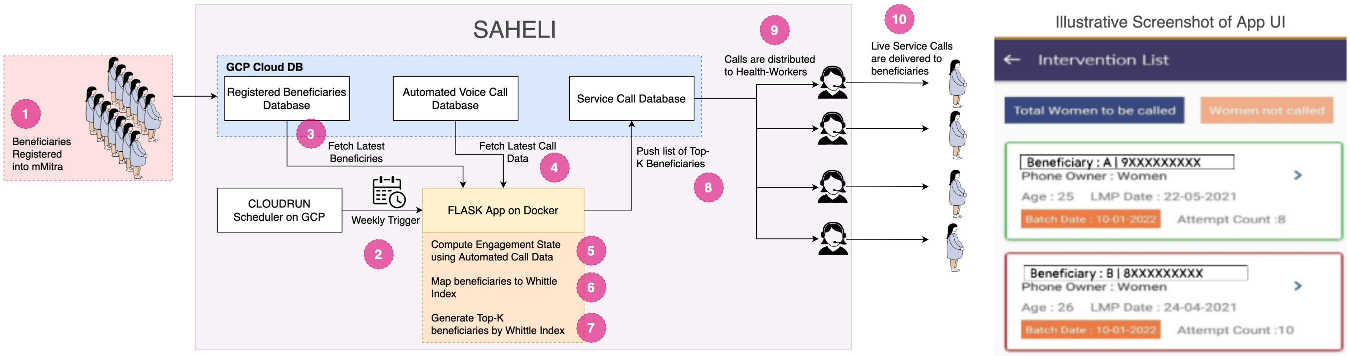

Figure 2: Pipeline of Deployed System. Beneficiary information on app UI is available only to the health worker in charge.

an average rate of $20K$ beneficiaries per month with a budget of 1000 calls per week.

SAHELI streamlines the entire deployment workflow in a singular pipeline, and automates its orchestration and execution, making this process computationally efficient, cost-effective, and easy to debug. As more beneficiaries get enrolled periodically, the beneficiary cohort in the application can now be updated automatically. In our AI pipeline we focused on identifying the right set of beneficiaries to call, and not on automating the contents of the service call. *This is a key design choice in* SAHELI: we thus complement the human-to-human engagement between the health worker and the beneficiary, and together they contribute towards aiding a particular beneficiary and driving higher engagement with the mHealth program. This model of working together with the health workers embodies ARMMAN's core 'tech plus touch' philosophy (ARMMAN 2008) and is essential to our successful outcomes.

Pipeline Description

This section describes the modules in the AI pipeline for both the offline model training and the online model execution. The offline model creation begins with the processing of the training data (i.e. historic data from past mHealth studies), clustering of processed data, and the RMAB modeling per cluster. The transition probabilities and the Whittle indexes are then learned per cluster. Additionally, a mapping from socio-demographic features of a beneficiary to a cluster is also learned offline. This mapping is used to treat a new beneficiary during model execution – transition probabilities and Whittle index values for the new beneficiary are given by the corresponding values of the beneficiary's mapped cluster. These individual modules are now described. For data privacy reasons, the data pipeline only uses anonymized data and no personally identifiable information (PII) is made available to the AI models.

Data Processing: We train the model on a dataset obtained from historic data collected by ARMMAN, consisting of demographic features and listenership patterns. However, during the pre-deployment trials, we observed some anomalous engagement behaviors – the engagement behavior for some beneficiaries was extremely spiky and unexpected. Figures 3(a) and (b), show two such anomalous groups with

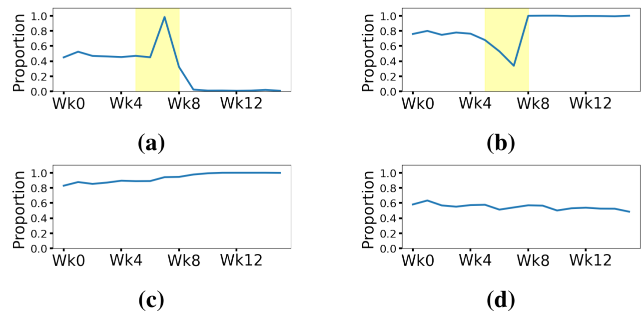

Figure 3: Figures (a) and (b) show anomalous engagement behavior while figures (c) and (d) are genuine behaviors. The y-axis shows the proportion of cluster-population in engaging state.

a clear peak and dip contrasted with groups having genuine engagement behavior. Upon investigation we found that this spiky behavior resulted from unanticipated real-world events like network outages.

We detect and exclude such anomalies from SAHELI's data training pipeline. We first group beneficiaries based on their passive transition probabilities. For grouped beneficiaries, we then obtain a running mean of their engagement over time where the mean is calculated over a window of 3 weeks. We filter out all groups with more than 20% change in running mean engagement within a week. Figures 3(c) and (d) show two groups that don't exhibit anomalous behavior and are maintained in the data pipeline.

Additionally, further discussions with ARMMAN pointed out long-term engagement issues in some beneficiaries, such as the registration of a wrong or out-of-service phone number, or the beneficiary not being pregnant. Live service calls in these cases are not productive. Thus, as a pre-processing step, we do not consider beneficiaries who have not listened to any automated voice calls in the past 6 weeks.

Clustering: We face a data scarcity and skew challenge in our domain. Specifically, our training dataset comprises beneficiaries from our own past studies where intervention data is available for only a limited set of these beneficiaries. Thus, to define the parameters of the RMAB model, we cluster beneficiaries as an effective way of addressing

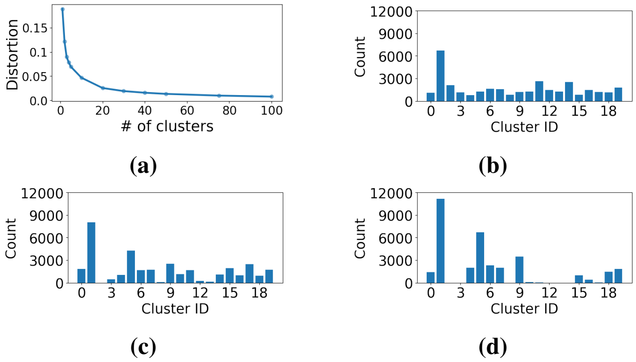

Figure 4: Figure (a) shows elbow plot with distortion for varying number of clusters. Figures (b), (c), and (d) show the distribution of predicted clusters using the Feature Only (FO), Feature and Warm-up (FW), and Warm-up Only (WO) mapping functions.

data scarcity. We cluster the beneficiaries per their transition behaviors for passive actions using *k-means* clustering. We obtain transition probabilities for each of these clusters by aggregating their transitions as a whole.

However, the optimal number of clusters is a design choice not readily addressed by *k-means*. We experimented with the number of clusters ranging from 1 to 100, and looked at the *distortion* metric. Distortion is the sum of squared distances of each point from its corresponding centroid, where smaller distortion implies better clustering. We plot the distortion values for multiple number of clusters and find 20 to be the ideal choice using elbow-method. The results are shown in Figure 4a where the x-axis is the number of clusters and the y-axis is the distortion value. This has the added advantage of offering computational frugality.

Mapping Features to Clusters: When a new beneficiary enrolls into the system, the system only knows about their demographic data. We therefore need to learn a mapping of a beneficiary's socio-demographic features to clusters, to enable inferring transition probabilities and Whittle indexes for newly enrolled beneficiaries (step 6 in Figure 2). We experimented with different mapping functions to identify the best one: Features Only (FO) mapping - beneficiaries' socio-demographic features only; Warm-up Only (WO) mapping - transition probabilities computed from warm-up period (first 6 weeks post enrollment); and lastly Feature and Warm-up (FW) mapping - using a combination of the above two.

We compute Mean Absolute Error between predicted and ground truth passive transition probabilities as a performance metric and found them as $[0.40, \mathbf{0.37}, 0.38]$ for FO, FW, and WO strategies respectively. In addition to MAE, we plot the distribution of beneficiaries predicted in different clusters (refer Figures 4(b), (c) and (d)). Having a sparse cluster distribution is undesirable since large clusters lowers the granularity of Whittle index planning. As an extreme example, if all beneficiaries are mapped to a single cluster, they would all have the same transition probability and thus the same Whittle indexes. Since the cluster size is now much larger than the number of arms to be pulled, the beneficiaries within that cluster would be chosen randomly for receiving service calls, which would degrade the performance.

Thus, to ensure equitable cluster distribution, we computed Entropy and Gini index values for the predicted distribution of number of beneficiaries per cluster. Entropy values came out to be $[\mathbf{2.81}, 2.56, 2.04]$ for FO, FW, and WO respectively, and Gini indexes were $[\mathbf{0.29}, 0.48, 0.57]$. Given the error similarities for the three strategies, and higher entropy / lower Gini index implies more equitable clusters, we chose FO as our strategy.

There are additional improvements for efficiency in SAHELI. We refer the reader to Verma et al. (2023) for further details.

Application Use and Payoff

We now discuss the impact of SAHELI on both the beneficiaries as well as the AI community in more detail. SAHELI is deployed and in continuous use at ARMMAN. It has already reached 100K beneficiaries, and is on track to reach one million beneficiaries by the end of 2023. We provide a summary of Impact from SAHELI in Table 2.

Engagement Results

In order to evaluate the impact of live service calls through SAHELI, we study the engagement behavior of a cohort of 5000 beneficiaries for 12 weeks, registered between February 2022 to April 2022. Additionally, we create a holdout set of beneficiaries registered in the same time period but are not given any live service calls (we obtained ethical approvals before our studies; see section Responsible AI practices for further discussion). We make sure that both the SAHELI and holdout groups have equal number of beneficiaries, equal number engaging beneficiaries at the start of experiment, and similar socio-demographic features.

Figure 5(a) shows how many engagements did not occur in the holdout group that occurred in the SAHELI group, aggregated cumulatively across months. It demonstrates that the SAHELI group received significant benefit with an additional 328 engagements over the holdout group cumulatively at the end of three months. We also measured the difference in terms of time spent listening to mMitra voice calls. More time spent implies more content exposure, as well as better adherence with the mHealth program. In particular, by the end of month 3, the SAHELI group had listened to $59,336$ seconds ($\sim$ 12 seconds per beneficiary, but please see analysis below) *more of content than the holdout group* (Figure 5(b)). Similar to Mate et al. (2022), we define the relative improvement in listenership metric over the holdout group as

$$\%\ improvement = \frac{\Delta\ listenership\ (SAHELI, holdout)}{listenership\ in\ holdout} \quad (1)$$

As the holdout group has 1075 drops in engagements and $127,711$ seconds drop in duration of calls listened to over three months, SAHELI prevented **drop in engagements by $\mathbf{30.5}\%$** with an **additional content exposure of $\mathbf{46.4}\%$** in comparison to the holdout group. This analysis demonstrates SAHELI's success in achieving our core objectives of improving information dissemination.

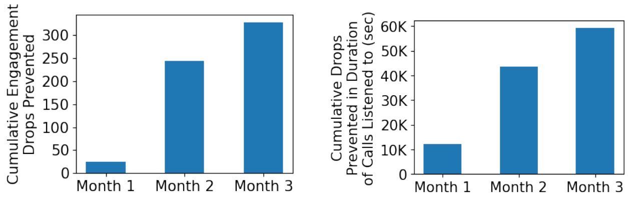

(a) **(b)**

Figure 5: (a) Prevention in drop in engagement (cumulative) (b) Increased time spent listening to calls (cumulative)

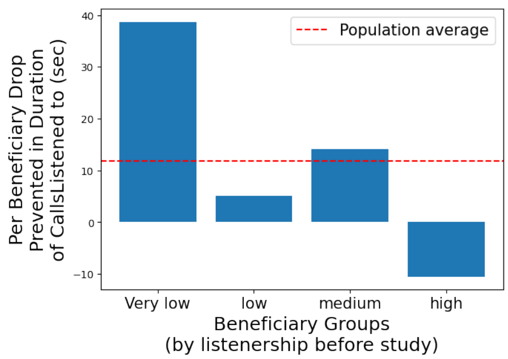

Figure 6: Increased time spent listening to calls (over three months). The metric is shown for beneficiaries belonging to *very low, low, medium and high* quartiles of listenership before the start of study.

Who is Benefitted from SAHELI?

In order to determine the characteristics of beneficiary who gain the most from SAHELI, we divide the 5000 beneficiaries in our cohort based on two criterion

1. Listenership prior to the start of study
2. Gestational age at the time of enrollment

First, we consider the listenership of beneficiaries one month prior to start of live service calls delivered through SAHELI. In this time period, we calculate the mean duration of calls listened to every week. Based on this metric, we divide the 5000 strong cohort into quartiles of listenership - *very low, low, medium and high*. These quartiles thus characterize the initial behaviour of beneficiaries. Next, we repeat the same steps for the holdout population which doesn't receive any service calls. Finally, we plot how many more seconds of mMitra content is listened by every beneficiary in the quartiles in SAHELI group as compared to the same quartiles in the holdout group (Figure 6).

While the population average increase in content listenership is $\sim$ 12 seconds, beneficiaries with different listenership profiles before being exposed to SAHELI show very distinctive behaviours. Specifically, the *very low* quartile of beneficiaries gain the most in SAHELI, with 39 seconds additional content listenership over the holdout group. In absolute terms, the *very low* quartile in holdout group has per beneficiary 30 seconds of increase in duration of calls listened to over three months while the SAHELI group has per beneficiary 69 seconds of increase in duration of calls listened in the same time period for the same quartile. Thus, using Equation 1, we note that in relative terms, the *very low* quartile has 130% additional content exposure in comparison to the holdout group.

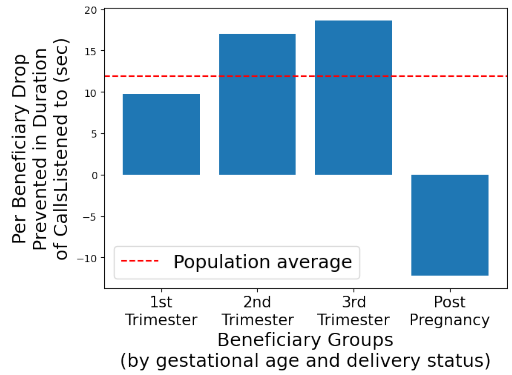

Figure 7: Increased time spent listening to calls (over three months). The metric is shown for pregnant mothers in their 1st, 2nd and 3rd trimesters of pregnancy and for beneficiaries who have already delivered.

Impact from SAHELI	
Beneficiairies served	100K
In continous use since	April 2022
Relative engagement drops prevented over holdout group*	30.5%
Additional average per beneficiary content exposure over holdout group*	12 seconds
Relative increase in content exposure over holdout group*	46.4%
For bottom 25 percentile of listeners, Additional average per beneficiary content exposure over holdout group*	39 seconds
For bottom 25 percentile of listeners, relative increase in content exposure over holdout group*	130%

Table 2: A summary of impact from SAHELI. * refers to results from a sample of 5000 beneficiaires.

For the second criterion, we consider the gestational age of beneficiaries and their delivery status at the time of enrollment. For pregnant women, we use the gestational age at the time of enrollment to calculate their pregnancy trimester. Similar to Figure 6, in Figure , we plot for every gestational age bucket, how many additional seconds of mMmitra calls are listened by every beneficiaries in the SAHELI group as compared to the holdout group. Specifically, we observe that beneficiaries close to the delivery date (higher trimester) have greater benefit from being in the SAHELI group.

Impact of Live Service Calls

We performed a qualitative study to understand the experiences and challenges faced by healthcare workers upon the introduction of SAHELI. We conducted a total of 24 interviews, 2 focus group discussions, and approximately 90 hours of observation over a period of six weeks. Conclusions were drawn by analyzing interview transcripts (audio recorded with consent). We found that with SAHELI, health workers were able to have more interactive conversations with their beneficiaries as they were aware that they had to provide support to people who were at high risk of drop off otherwise. In one of the interviews, one of the health workers mentioned that:

Women don't remember that they registered by the time they go home. A lot is happening during their visit. When they get a call, then they remember. We are able to do this better now since we know we are targeting those who need this call the most.

We also investigated the reasons for why live service calls helped improve engagement with ARMMAN's mMitra mHealth program from the perspective of the beneficiary. Specifically, we conducted a follow-up study with a sample of beneficiaries who were given live service calls one year ago. We could successfully reach out to 306 beneficiaries, out of which 134 recalled the details of the service call from a year ago. Table 3 shows the responses to our follow-up study by these 134 beneficiaries. Particularly, 50.75% beneficiaries engaged more with mMitra calls after getting more information about the program. The service calls also helped improve listenership by making logistical updates such as updating delivery date (9.7%), changing time slot of receiving the call (8.21%) or updating the phone number (2.99%).

Did the call help you to listen to the mMitra calls more regularly?	**# of Beneficiaries**	**% of Beneficiaries**
Yes, after getting more information about mMitra, I am listening to the calls more regularly	68 (in 134)	50.75%
Not really	30	22.39%
Yes, after updating my delivery date, I was able to get the right information	13	9.7%
Yes, after changing time slot, I am able to listen to the calls more regularly	11	8.21%
Have not asked my wife	4	2.99%
Yes, after changing the number, I am able to listen to the calls more regularly	4	2.99%
Any other	4	2.99%

Table 3: Follow-up study responses

Lessons Learned

Over the course of one year of our experiments moving from Pilot study to Deployment, we learned several lessons along the way. Most importantly, we learned that even a successful pilot study can't be translated as-is into a full-scale deployment, and that several considerations are critical for wide-scale adoption of AI tools and scaling up of impact.

Selecting the right problem: There are multitude of problems that require to be solved to address the needs of the underserved communities. In our interactions with ARMMAN, *we realized that we could create the most impact with our techniques by improving the selection of the right beneficiaries for manual intervention*, as opposed to automating the communication with the beneficiary. Our choice of problem is consistent with the 'tech plus touch' philosophy of ARMMAN (2008), and ensures that we complement the human expertise of the health worker. This way, each chosen beneficiary continued to have a one-on-one interaction with a health worker, while simultaneously improving the overall engagement with the mHealth program.

Immersion into the real-world problem: We learned that immersing in the working of a NGO and public health infrastructure is critical in understanding the context of the problem. The authors went on multiple field visits to understand the stakeholders involved in the mMitra's workflow. The health workers interact with the beneficiaries across multiple mHealth programs, and thus can speak to the needs and behaviors of the beneficiaries. For instance, upon interacting with these health workers, we understood how telecom outages lead to more anomalous and incomplete data than we had anticipated. *These field visits forced us to reevaluate our assumptions, and led to better data processing and modeling choices*, as discussed in the earlier sections. For instance, after these discussions, we incorporated a new anomaly detection mechanism in our data pipeline.

End-to-end integration testing: We also ran into several issues in our end-to-end integrated pipeline. On one occasion, we saw poor results because the data schema had evolved in the data storage pipeline at ARMMAN. *Testing of our application required our NGO partner to be equally involved in the validation of* SAHELI*'s outputs – as domain experts, they are better equipped to identify counter-intuitive behaviors.* Our experiences uncovering issues in the end-to-end pipeline led to improved communication practices, better documentation and tighter test goals. Social good applications like SAHELI have real-world consequences for beneficiaries in underserved communities, and it is critical that there be a real partnership for testing and integration.

Conclusion

In this paper, we presented SAHELI, the first ever deployment of RMABs in the public health domain for allocation of limited resources. SAHELI is built on an improved and robust framework that both predicts RMAB parameters and computes optimal policies for it, in contrast with most past research that has only focused on computing optimal policies. It has been built with careful design choices inspired by close interactions with all stakeholders. It incorporates numerous lessons learned by embedding ourselves in the real-world domain. SAHELI has been deployed on cloud infrastructure with an emphasis on frugality, and has reached out to 100K beneficiaries so far and aims to reach 1 million by 2023. Furthermore, we also discuss the importance of responsible AI practices in deploying AI systems at scale, especially in the social domain. This work serves as an important case study for AI researchers and NGO's alike to take ML models from the lab and deploy them in the field.

References

ARMMAN. 2008. About ARMMAN. https://armman.org/about-us. Accessed: 2022-08-12.

ARMMAN. 2019. Assessing the Impact of Mobile-based Intervention on Health Literacy among Pregnant Women in

Urban India. https://armman.org/wp-content/uploads/2019/09/Sion-Study-Abstract.pdf. Accessed: 2022-08-12.

Corotto, P. S.; McCarey, M. M.; Adams, S.; Khazanie, P.; and Whellan, D. J. 2013. Heart failure patient adherence: epidemiology, cause, and treatment. *Heart failure clinics*, 9(1): 49–58.

Eysenbach, G. 2005. The Law of Attrition. *J Med Internet Res*, 7(1): e11.

Glazebrook, K.; Ruiz-Hernandex, D.; and Kirkbride, C. 2006. Some indexable families of restless bandit problems. *Adv. Appl. Probab*, 643–672.

Gupta, K.; Roy, S.; Poonia, R.; Nayak, S. R.; Kumar, R.; Alzahrani, K.; Alnefaie, M.; and Al-Wesabi, F. 2022. Evaluating the Usability of mHealth Applications on Type 2 Diabetes Mellitus Using Various MCDM Models. *Healthcare*, 10: 4.

Jakob, R.; Harperink, S.; Rudolf, A. M.; Fleisch, E.; Haug, S.; Mair, J. L.; Salamanca-Sanabria, A.; and Kowatsch, T. 2022. Factors Influencing Adherence to mHealth Apps for Prevention or Management of Noncommunicable Diseases: Systematic Review. *J Med Internet Res*, 24(5): e35371.

Killian, J. A.; Wilder, B.; Sharma, A.; Choudhary, V.; Dilkina, B.; and Tambe, M. 2019. Learning to Prescribe Interventions for Tuberculosis Patients Using Digital Adherence Data. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.

Liao, P.; Greenewald, K.; Klasnja, P.; and Murphy, S. 2020. Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(1): 1–22.

Liu, K.; and Zhao, Q. 2010. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory*, 56(11): 5547–5567.

Mate, A.; Madaan, L.; Taneja, A.; Madhiwalla, N.; Verma, S.; Singh, G.; Hegde, A.; Varakantham, P.; and Tambe, M. 2022. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11): 12017–12025.

Meh, C.; Sharma, A.; Ram, U.; Fadel, S.; Correa, N.; Snelgrove, J.; Shah, P.; Begum, R.; Shah, M.; Hana, T.; Fu, H.; Raveendran, L.; Mishra, B.; and Jha, P. 2022. Trends in maternal mortality in India over two decades in nationally-representative surveys. *BJOG: An International Journal of Obstetrics & Gynaecology*, 129.

Meyerowitz-Katz, G.; Ravi, S.; Arnolda, L.; Feng, X.; Maberly, G.; and Astell-Burt, T. 2020. Rates of Attrition and Dropout in App-Based Interventions for Chronic Disease: Systematic Review and Meta-Analysis. *J Med Internet Res*, 22(9): e20283.

Papadimitriou, C. H.; and Tsitsiklis, J. N. 1999. The complexity of optimal queuing network control. *Mathematics of Operations Research*, 24(2): 293–305.

Pilote, L.; Tulsky, J. P.; Zolopa, A. R.; Hahn, J. A.; Schecter, G. F.; and Moss, A. R. 1996. Tuberculosis Prophylaxis in the Homeless: A Trial to Improve Adherence to Referral. *Archives of Internal Medicine*, 156(2): 161–165.

Pollack, M. E.; Brown, L.; Colbry, D.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; Engberg, S.; Matthews, J. T.; Dunbar-Jacob, J.; McCarthy, C. E.; et al. 2002. Pearl: A mobile robotic assistant for the elderly. In *AAAI workshop on automation as eldercare*, volume 2002. AAAI, 2002, Edmonton, Alberta, Canada.

Puterman, M. L. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.

Qian, Y.; Zhang, C.; Krishnamachari, B.; and Tambe, M. 2016. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In *AAMAS*.

Tshikomana, R. S.; and Ramukumba, M. M. 2022. Implementation of mHealth applications in community-based health care: Insights from Ward-Based Outreach Teams in South Africa. *PLOS ONE*, 17(1): 1–15.

Tuldrà, A.; Ferrer, M. J.; Fumaz, C. R.; Bayés, R.; Paredes, R.; Burger, D. M.; and Clotet, B. 1999. Monitoring Adherence to HIV Therapy. *Archives of Internal Medicine*, 159(12): 1376–1377.

Verma, S.; Singh, G.; Mate, A. S.; Verma, P.; Gorantala, S.; Madhiwalla, N.; Hegde, A.; Thakkar, D. H.; Jain, M.; Tambe, M. S.; and Taneja, A. 2023. Deployed SAHELI: Field Optimization of Intelligent RMAB for Maternal and Child Care. In *Innovative Applications of Artificial Intelligence (IAAI)*.

Weber, R. R.; and Weiss, G. 1990. On an Index Policy for Restless Bandits. *Journal of Applied Probability*, 27(3): 637–648.

Whittle, P. 1988. Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.*, 25(A): 287–298.

World Bank, . 2020. *Poverty and shared prosperity 2020: Reversals of fortune*. The World Bank.

World Health Organization (WHO). 2020. Newborns: improving survival and well-being. https://www.who.int/news-room/fact-sheets/detail/newborns-reducing-mortality#:~:text=Neonates,in\%20child\%20survival\%20since\%201990. Accessed: 2022-08-12.

Appendix

Ethics

We recognize the responsibility associated with deploying real-world AI systems that impacts underserved communities. In our approach, we have iteratively designed, developed and deployed the system in constant coordination with an interdisciplinary team of ARMMAN's field staff, social work researchers, public health researchers and ethical experts. Particularly, all experiments, field tests and the deployment were performed after obtaining approval from ethics review board at both ARMMAN and Google.

Consent and Data Usage

The consent for participating in the mMitra program is received from beneficiaries in both written form at the time of registration and digitally via a missed call. Additionally, all the data collected through the program is owned by the NGO and only the NGO is allowed to share data. This dataset will never be used by Google for any commercial purposes. SAHELI's data pipeline only uses anonymized data and no personally identifiable information (PII) is made available to the AI models. The data exchange and use was thus regulated through clearly defined exchange protocols including anonymization, read-access only to researchers, restricted use of the data for research purposes only, and approval by ARMMAN's ethics review committee.

Universal Accessibility of Health Information

To allay further concerns: SAHELI focuses on improving quality of service calls and does not alter, for any beneficiary, the accessibility of health information. All participants will receive the same weekly health information by automated message regardless of whether they are scheduled to receive service calls or not. The service call program does not withhold any information from the participants nor conduct any experimentation on the health information. The health information is always available to all participants, and participants can always request service calls via a free missed call service.