Expanding Impact of Mobile Health Programs: SAHELI for Maternal and Child Care

Shresth Verma^{1*}, Gargi Singh^{1*}, Aditya Mate^{12†}, Paritosh Verma^{13†}, Sruthi Gorantla^{14†}, Neha Madhiwalla⁵, Aparna Hegde⁵, Divy Thakkar¹, Manish Jain¹, Milind Tambe¹, Aparna Taneja¹

¹Google Research India ²Harvard University ³Purdue University ⁴Indian Institute of Science, Bangalore ⁵ARMMAN

{vermashresth, gargisingh}@google.com, aditya_mate@g.harvard.edu, verma136@cs.purdue.edu, gorantlas@iisc.ac.in, {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. Nongovernmental 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 Multiarmed 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 $\sim 130K$ 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.

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.

- 10	States Covered in	19
	India	
	Partner NGOs	40
	Partner Hospitals	97
	Health Workers Trained	235K
	Beneficiaries	27.2M
ZADA	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 due to 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 deployed Restless Multi-Armed Bandits (RMAB) model for public health – together with our NGO partner ARMMAN (ARMMAN 2008) to help improve the quality of service of their mHealth program focusing on maternal and child care. India suffers from high maternal and neonatal mortality rates (Meh et al.

^{*}Joint first author

[†]Work done during an internship at Google Research India Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2022; World Health Organization (WHO) 2020), and AR-MMAN (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 interdisciplinary 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 Responsible AI practices to address ethical considerations for deploying an AI system for impact in underserved communities, particularly in this non-western context.

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 Github¹.

Related Work

While several works in the healthcare domain have studied patient adherence for diseases like HIV (Tuldrà et al. 1999), cardiac problems (Son et al. 2010; 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); Brisimi et al. (2018) 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), multichannel 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.

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 (~ 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 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 dropout 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 beneficiaries on a weekly basis to encourage beneficiaries' participation, address requests/ complaints, and attempt to prevent engagement drops. This limitation in number of service calls raises the key question of deciding which beneficiaries to pick for live service call 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 {S, A, r, P} consists of a set of states S, a set of actions A, a reward function $r : S \times A \times S \mapsto \mathbb{R}$, and a transition function P, where $P_{s,s'}^{\alpha}$ is the probability of transitioning from state s to s' when action α is chosen. The reward function in our set up is given as $r(s, \alpha, s') = s'$. An MDP policy $\pi : S \mapsto A$ maps to the choice of action to take at each state. The long-

¹https://github.com/armman-projects/SAHELI

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 $S = \{0, 1\}$, 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 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, i.e., deliver live service calls, so as to maximize the total reward which is the 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.

The pilot results demonstrated that the RMAB method cuts ~ 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 (see (Mate et al. 2022) for more details). 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

Improvements	RMAB over CSOC	RMAB over RR	RR over CSOC
% reduction in to- tal beneficiary en- gagement drops	32.0%	28.3%	5.2%
p-value estimate	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 showed by Mate et al. (2022).

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' gestation age are sent with the help of a telecommunication provider (step 4). The metadata 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.

The AI pipeline (described in the next section) for a dynamically growing population is deployed on infrastructure hosted on Google Cloud Platform (GCP). The AI pipeline is wrapped as an application using Flask, which is containerized using Docker. The docker image is created to contain the requisite code scripts for the AI pipeline with apt environment requirements. Our default GCP container settings are to use 6 vCPUs and 16GiB memory. A weekly sched-



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

uler job on GCP triggers the Flask application, which then generates the list of beneficiaries.

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 health worker sees a list of beneficiaries that he/she can call, along with certain features like number of call attempts. They can also click on a particular beneficiary and see more information about the beneficiary and past calls with them (not shown). The calls are made through the week 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 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, costeffective, and easy to debug. As more beneficiaries get enrolled periodically, the beneficiary cohort in the application can now be updated automatically.

Health workers can then make the calls (step 10 in Figure 2) to these beneficiaries motivating them to listen to the voice messages and address any logistic issues (e.g. time slots, language of communication, and others) that might be affecting their engagement. As we show later in the paper, motivating the beneficiaries is key to driving adherence. However, it bears repeating that given the limited availability of the health workers, they can only make a limited number of calls. 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 (AR-MMAN 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), shows two such anomalous groups with 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



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.

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 of 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 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 clustering approach 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' sociodemographic 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, 0.37, 0.38] for FO,



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.

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 [2.81, 2.56, 2.04] for FO, FW, and WO respectively, and Gini indexes were [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.



Figure 5: Index computation is significantly faster with the infinite sleeping approximation.

RMAB Modeling and Whittle Index Computation: These transition probabilities per cluster are used to compute Whittle indexes for all beneficiaries, similar to Mate et al. (2022), i.e., computing $2 \times k$ unique indexes where k is the number of clusters. There are two Whittle indexes per cluster as beneficiaries may be in the engaging or nonengaging states. Whittle index indicates the benefit of performing an active action on a beneficiary: higher Whittle indexes are chosen to receive service calls (step 7 in Figure 2). By mapping beneficiaries to clusters, the Whittle indexes can be pre-computed per cluster at the beginning of the deployment, thus providing a frugal solution ideal for large scale deployment with minimal resources.

Frequency of Repeated Live Service Calls: We initially enforced a frequency restriction that required ensuring no beneficiary be called more than once in η +1 weeks (we set $\eta = 3$). Algorithmically, we implement this by appending η sets of dummy 'sleeping states' to the state space that we force the beneficiaries to transition through each time they are called. This augmentation yields a state space of size $2\eta + 2$ and a transition matrix of size $(2\eta + 2) \times (2\eta + 2)$. However, our pilot tests reveal that repeat calls made within just $\eta = 3$ weeks are less effective. For instance, we observed that 30% of 'Non-engaging' beneficiaries converted to 'Engaging' due to the first service call; however this number drops to 20% for repeat calls made just three weeks later. To address this issue, along with the subject matter experts at ARMMAN, we increased the sleeping period, η , to 12 weeks.

Frugality of System Design

Successful deployments of AI systems like SAHELI in social good settings requires conscious focus on frugality across the system design. This frugality means reducing both the direct costs (e.g. number of calls) and indirect costs (e.g. computational requirements) on our NGO partners. Here are some design choices in SAHELI that have led to frugality in its operations:

1. Clustering of beneficiaries allows us to compute transition probabilities and Whittle indexes at a cluster level as opposed at the beneficiary level. Since we use 20 clusters for thousands of beneficiaries, it provides a significant scaleup in performance, while simultaneously reducing data demands for learning RMAB model parameters.

2. As described above, we updated the 'sleeping states' parameter η to 12. However, this increases the Whittle index computation time sharply, owing to a bulky transition matrix of size 26×26 . With frugality in mind, we use the insight that a sleeping constraint with large η can be approximated as a permanent sleeping constraint, akin to setting η to $+\infty$, for the purposes of index computation. This is because in index computation, the contribution of reward terms appearing after η timesteps is discounted by a factor of γ^{η} ($\gamma < 1$), which precipitously diminishes to zero. This simplification compresses the transition matrix to 4×4 , and unlocks a $25 \times$ speedup in index computation, as shown in Figure 5.

3. Lastly, multiple frugal design choices were made in the orchestration of cloud infrastructure. Specifically, we run our services on-demand using a task scheduler on default container settings of 6 vCPUs and 16GiB memory.

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 130K 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 6(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 6(b)). Similar to Mate et al. (2022), we define the relative improvement in listenership metric over the holdout group as

% improvement =
$$\frac{\Delta \text{ listenership (SAHELI, holdout)}}{\text{ listenership in holdout}}$$
(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 30.5**% with an **additional content exposure of 46.4**% in comparison to the holdout group. This analysis demonstrates SAHELI's success in achieving our core objectives of improving information dissemination.



Figure 6: (a) Prevention in drop in engagement (cumulative) (b) Increased time spent listening to calls (cumulative)

Who is Benefitted from SAHELI?

In order to determine the characteristics of beneficiary who gain the most from SAHELI, we divide the 5000 beneficia-



Figure 7: 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.

ries 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 7).

While the population average increase in content listenership is ~ 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.

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 7, 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.



Figure 8: 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	130K	
In continous use since	April 2022	
Relative engagement drops prevented over hold- out group*	30.5%	
Additional average per beneficiary content exposure over holdout group*	12 seconds	
Relative increase in content exposure over hold- out 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 in- crease in content exposure over holdout group*	130%	

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

Impact of Live Service Calls

We performed a qualitative study to understand human-AI collaboration due to the AI system. We conducted a total of 24 interviews, 2 focus group discussions, and approximately 90 hours of observation. We found that healthcare workers engaged positively with targeted predictions through the AI system that integrated into their day-day workflows seam-lessly. It helped them improve the engagement of beneficiaries, provided an opportunity to support them in their care journeys and understand their needs.

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 Benefi- ciaries	% of Benefi- ciaries
Yes, after getting more information about mMi- tra, 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

Fairness of the RMAB model

Model fairness in non-western contexts has not received much attention in the literature (Sambasivan et al. 2021). Responsible AI principles of the Government of India's NITI AAYOG (2021) for example, requires non-discrimination based on sensitive markers like caste and religion. These sensitive data are specifically not collected by ARMMAN for mMitra, thereby, making it inaccessible to SAHELI's AI models. We worked with public health and field experts to evaluate other indicators such as education, and income levels that signify markers of socio-economic marginalization. ARMMAN's goals for SAHELI are to favor beneficiaries of lower income and lower education levels for service calls. We conducted a post-hoc analysis of the deployment to evaluate if SAHELI indeed met such preferences.

Figure 9(a) shows the distribution of beneficiaries aggregated across SAHELI's enrollments split into different education levels in India. We compare those who were chosen for live service calls by SAHELI versus the enrolled



Figure 9: Distribution of (a) education (highest education received) and (b) income (monthly family income in Indian Rupees) across cohort that received service call and the whole population.

population. The x-axis portrays the education levels; for instance grade 1-5 represents primary school, grade 6-9 middle school, 10th pass junior high, and 12th pass represents senior high school. The y-axis is the % of beneficiaries per education category. For instance, SAHELI calls 5.5% of beneficiaries who had very little formal education or were illiterate, whereas this group was 2.8% of the overall enrolled population.

We did a similar analysis split by income as depicted in Figure 9(b). The x-axis contains buckets of average monthly income of the beneficiary household in Indian Rupees, and the y-axis denotes the % of beneficiaries in that income category. As an example, the category '5K-10K' contains around 30% of the beneficiaries in the population, and almost 40% of the beneficiaries who received a service call.

Both these figures show that SAHELI favors the beneficiaries in the 'illiterate' education category and in the '5K-10K' income category. This distribution is in line with AR-MMAN's goals – SAHELI favors beneficiaries of lower income and lower education levels for service calls.

Enabling New Research

From identifying the right problem to solve, to creating an AI solution, testing it in pilot, iterating on learnings and finally, establishing an end-to-end integrated system, we made our journey to this deployment. With this work, we provide other AI researchers an important case study to take an AI model from the lab out on the field. In our pursuit of deployment of SAHELI, we uncover several research challenges, e.g., we overcame the challenges of data scarcity and frugal design. Our work hopefully inspires additional research in robust and computationally efficient approaches for RMABs and other AI applications for mHealth. Further, we are improving SAHELI by incorporating recent advances in Decision Focused Learning for RMAB problems (Wang et al. 2023) which opens up new direction of real-world large-scale applications of Decision Focused Learning.

Responsible AI Practices

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 comprised of ARMMAN's field staff, social work researchers, public health researchers and ethical experts. Along with seeking ethical approvals through review boards at Google and ARMMAN, we have taken additional steps to constantly monitor and mitigate the risks associated with SAHELI by abiding with AI principles at Google (2018) as well as key policy making bodies in India such as the NITI AAYOG (2021). Our success draws attention to the practices around responsible AI including ethics, fairness and accountability in the non-western context (Sambasivan et al. 2021) where SAHELI is deployed. We now discuss three of the core Responsible AI principles that impacted the design of SAHELI.

Socially beneficial: The intent of this work is to bring the power of AI in service to some of the most marginalized communities in the global south. The challenges faced by our team were limited resources in every dimension – limited data on the beneficiaries, limited compute available to the NGO, and limited health workers to make the outreach calls. Thus, we had to develop new algorithms that were not data hungry, and were bounded in their computational requirements. To that affect, SAHELI is the first large-scale deployment of RMABs for public health.

Avoid reinforcing unfair bias: As discussed in the previous section, we have undertaken extensive analysis to study model's fair treatment of beneficiaries.

Incorporate privacy design principles: We take significant measures to ensure participant consent is understood and recorded in a language of the community's choice at each stage of the program. Data stewardship resides in the hands of the NGO, and only the NGO is allowed to share data. This dataset will never be used by Google for any commercial purposes. In this dataset, sensitive features such as caste and religion are never collected and stored. SAHELI's data pipeline only uses anonymized data and no personally identifiable information (PII) is made available to the AI models. Lastly, domain experts at ARMMAN have been deeply involved in the development and testing of SAHELI and have provided continuous input and oversight in data interpretation, data consumption and model design.

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 in to a full-scale deployment, and that several considerations are critical for widescale 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 ARM-MAN, we realized that we could create the most impact with our techniques by improving the selection of the right ben-

eficiaries 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. We also understood the decreased value in utility of calling the same beneficiary again shortly after a previous call. 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, and impacted our choice of horizon (η) in our RMAB model.

Fairness of AI models: AI algorithms and datasets can reflect, reinforce, or reduce unfair biases. It is imperative on AI designers to seek to avoid unfair impacts on people, particularly on underserved and marginalized communities. As discussed in the section on Responsible AI practices, we worked with public health and field experts to demonstrate fairness of our approach. As we mentioned before, 94% of our potential beneficiary population are below WHO's poverty index. *Studying multiple socio-demographic attributes was essential to evaluate fairness of our approach.* We worked closely with ethics experts, the ARMMAN's ethics team, and Google's ethics teams and extensively validated the fairness of our models.

End-to-end integration testing: In addition to the lessons learned on problem selection and model development, 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 has 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 restless multi armed bandits 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 130k beneficiaries so far and aims to reach 1 million by 2023. Furthermore, in this paper, 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 communities alike to take ML models from the lab and deploy them in the field.

Acknowledgments

We thank Lovish Madaan and Aakriti Kumar for their insights and feedback through the project, Harshit Khaitan, Anirudh GP, and Hasan Rafiq for their support in deployment on GCP, and Manoj Sandur Karnick for aid with program management. We are also very grateful to Suresh Chaudhary, Sonali Nandalaskar and the entire mMitra team from ARMMAN for their continuous support since pilot in making this journey to deployment possible. Conflict of Interest: The authors have no conflicts of interest to report.

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.

Brisimi, T. S.; Chen, R.; Mela, T.; Olshevsky, A.; Paschalidis, I. C.; and Shi, W. 2018. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics*, 112: 59–67.

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.

Google. 2018. Artificial Intelligence at Google: Our Principles. https://ai.google/principles. Accessed: 2022-08-12.

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. *In Proceedings of AAAI*, 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 nationallyrepresentative 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.

NITI AAYOG. 2021. Responsible AI: Approach Document for India. https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf. Accessed: 2022-08-12.

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*.

Sambasivan, N.; Arnesen, E.; Hutchinson, B.; Doshi, T.; and Prabhakaran, V. 2021. Re-imagining Algorithmic Fairness in India and Beyond. *CoRR*, abs/2101.09995.

Son, Y.-J.; Kim, H.-G.; Kim, E.-H.; Choi, S.; and Lee, S.-K. 2010. Application of support vector machine for prediction of medication adherence in heart failure patients. *Healthcare informatics research*, 16(4): 253–259.

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.

Wang, K.; Verma, S.; Mate, A.; Shah, S.; Taneja, A.; Madhiwalla, N.; Hegde, A.; and Tambe, M. 2023. Scalable Decision-Focused Learning in Restless Multi-Armed Bandits with Application to Maternal and Child Health. arXiv:2202.00916.

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.

Author Biographies

Shresth Verma is a Pre-doctoral researcher at Google Research India in the Multi-Agent Systems for Social Impact group and an incoming PhD student at Harvard University Center for Research on Computation and Society. Previously, he was a Data Scientist at United Health Group where he developed healthcare risk models for millions of beneficiaries. His research interests include bandits, reinforcement learning and multi-agent systems with applications in public health and conservation.

Gargi Singh is currently a Pre-doctoral Researcher at Google Research India in the AI for Social Good team. She has a Bachelors and Masters in Computer Science and Engineering from the Indian Institute of Technology, Kanpur, India. She completed her masters thesis under the supervision of Dr. Ashutosh Modi in textual fine-grained emotion prediction. She received the best student paper award in ACII 2021 and the innovative application award in IAAI 2023. Her research interests encompass different areas in applied artificial intelligence.

Aditya Mate is an incoming Applied Scientist at Microsoft New England. He has a Ph.D. in Computer Science from Harvard University advised by Prof. Milind Tambe. Previously, he completed his B.Tech and M.Tech from the Indian Institute of Technology, Bombay (IIT Bombay). Aditya's research explores techniques from Sequential Decision Making, Machine Learning and Causal Inference for solving impactful real-world problems, especially in public health. His work on improving maternal healthcare has been recognized by the "IAAI 2023 Innovative Application Award" and he also won the second prize for his talk at the "INFORM 2022 Doing Good With Good OR Competition". His work has won the best paper awards at NeurIPS 2021 and NeurIPS 2020 workshops on machine learning for public health. His work has also been featured in the Nature Asia newsletter, Harvard newsletter and has appeared on the Google AI blog. Aditya is also fortunate to have previously been a research intern at IBM T.J. Watson Research Center and a research intern at Google Research (thrice).

Paritosh Verma is a PhD student in the Department of Computer Science at Purdue University. He is advised by Prof. Alex Psomas. His research focuses on problems at the interface of computer science and mathematical economics, particularly on algorithmic fairness in resource allocation settings and mechanism design.

Sruthi Gorantala is a PhD student at the Indian Institute of Science Bangalore. Her research interests are algorithms, optimization, and fairness. Her research is supported by a Google PhD Fellowship.

Neha Madhiwalla is Research Director at ARMMAN. She has a PhD in Social Sciences and has been involved in public health research for over 20 years. Her areas of expertise include maternal and child health, health systems, medical education and training and bioethics. She is presently involved in research on mhealth, exploring the use to telemedicine for the rehabilitation of children with malnutrition; humancomputer interactions in mhealth programmes for maternal and child health; using AI to improve the effectiveness of large-scale digital health communication programmes. She has worked at the Centre for Enquiry in Health and Allied Themes (CEHAT) and Centre for Studies in Ethics and Rights, both affiliated to the Anusandhan Trust, between 1987 and 2013. She worked as a Research Scientist at the Ramalingaswami Centre for Equity and Social Determinants of Health, Bangalore in 2016 and 2017. She has been a member of the research ethics committee at BARC Hospital, Mumbai since 2016. She has had a parallel career in development practice. She founded Chehak Trust, an organisation which has been involved in developing community-based models for education, adolescent health and disability rehabilitation. Chehak Trust was awarded the Edelweiss-Kubera Award for Innovation in Education in 2010. She received the MacArthur Fund for Leadership Fellowship (2002-2005) to develop a comprehensive adolescent sexual and reproductive health programme for adolescent girls.

Dr. Aparna Hegde is a Stanford and Cleveland Clinic trained, internationally renowned Urogynecologist. She is the founder of NGO ARMMAN, which creates scalable mHealth-based programs to impact maternal and child health in 20 states of India (reach: 36 million women and their children and 318,000 health workers). Dr. Hegde is Associate Professor (Hon) and head of the upcoming Department of Urogynecology at Cama Hospital, Mumbai, one of

India's first University-based Centers of Excellence in the field with a fellowship program. Dr. Hegde is an accomplished researcher, an NIH grantee, and a pioneer in 2D/3D pelvic floor ultrasound. She is the Chair of FIUGA, the foundation arm of IUGA (International Urogynecology Association), and the Chair of the IUGA Publications Committee. Dr. Hegde is a Skoll awardee (2020), senior Ashoka Fellow (2021) and was listed by Fortune Magazine as one of the World's 50 Greatest Leaders in 2021. She has also won the Women Transforming India award by Niti Aayog in 2022.

Divy Thakkar is Staff Program Manager at Google Research where he leads Strategy and Operations for several critical functions, He leads Google's academic research strategy and investments in India, Southeast and South Asia. Divy is an HCI researcher with interests in examining Human-AI interactions in low-resource and/or high–stakes environments. He co-leads efforts towards shaping a worldclass AI for Social Good team impacting lives at-scale.

Manish Jain is a member of Google Research India where he is leading an effort on understand, model and leverage propagation networks of information. He has also contributed towards research on digital health using reinforcement learning and decision theory. A few years prior to joining Google, he was the founding CTO of Avata Intelligence. He managed the startup for 8 years before it was acquired by Procore Technologies in early 2020. At the time of its acquisition, Avata's core product was a question-answering chatbot over corporate documents (built for employees of the corporation). He also holds a PhD in Computer Science, focusing on large-scale optimization, game theory and graphbased algorithms.

Milind Tambe is Principal Scientist and Director for "AI for Social Good" at Google Research; Concurrently, he is Gordon McKay Professor of Computer Science and Director of Center for Research in Computation and Society at Harvard University. He is recipient of the IJCAI John McCarthy Award, AAAI Feigenbaum Prize, AAAI Robert S. Engelmore Memorial Lecture Award, AAMAS ACM Autonomous Agents Research Award, INFORMS Wagner prize for excellence in Operations Research practice and MORS Rist Prize. He is a fellow of AAAI and ACM. For his work on AI and public safety, he has received Columbus Fellowship Foundation Homeland security award and commendations and certificates of appreciation from the US Coast Guard, the Federal Air Marshals Service and airport police at the city of Los Angeles.

Aparna Taneja is a researcher at the Multi Agents Systems for Social Impact team in Google Research India. She collaborates with several NGO's and academic partners in the fields of public health and conservation, and her primary focus is the collaboration with ARMMAN, an NGO focused on improving maternal and child health outcomes in India. Aparna joined Google in 2015 and worked with the Google Maps team in Zurich to improve search quality for Maps. She received her PhD in Computer Science at ETH Zurich under the supervision of Prof. Marc Pollefeys. She then pursued a postdoc at Disney Research Zurich. She received her Bachelors and Masters in Computer Science from IIT Delhi in 2006.