

# Optimizing Vital Sign Monitoring in Resource-Constrained Maternal Care: An RL-Based Restless Bandit Approach

Niclas Boehmer<sup>1,2</sup>, Yunfan Zhao<sup>1,3</sup>, Guojun Xiong<sup>1</sup>, Paula Rodriguez-Diaz<sup>1</sup>,  
Paola Del Cueto Cibrian<sup>4</sup>, Joseph Ngonzi<sup>5</sup>, Adeline Boatin<sup>4</sup>, Milind Tambe<sup>1</sup>

<sup>1</sup>School of Engineering and Applied Sciences, Harvard University, USA

<sup>2</sup>Hasso Plattner Institute, University of Potsdam, Germany

<sup>3</sup>GE Healthcare, USA

<sup>4</sup>Department of Obstetrics and Gynecology, Massachusetts General Hospital, Harvard Medical School, USA

<sup>5</sup>Mbarara University of Science and Technology, Uganda

niclas.boehmer@hpi.de, yunfanzhao@fas.harvard.edu, gjxiong@seas.harvard.edu, prodriguezdiaz@g.harvard.edu,  
pdelcueto@mgh.harvard.edu, jngonzi@must.ac.ug, adeline\_boatin@mgh.harvard.edu, milind\_tambe@harvard.edu

## Abstract

Maternal mortality remains a significant global public health challenge. One promising approach to reducing maternal deaths occurring during facility-based childbirth is through early warning systems, which require the consistent monitoring of mothers' vital signs after giving birth. Wireless vital sign monitoring devices offer a labor-efficient solution for continuous monitoring, but their scarcity raises the critical question of how to allocate them most effectively. We devise an allocation algorithm for this problem by modeling it as a variant of the popular Restless Multi-Armed Bandit (RMAB) paradigm. In doing so, we identify and address novel, previously unstudied constraints unique to this domain, which render previous approaches for RMABs unsuitable and significantly increase the complexity of the learning and planning problem. To overcome these challenges, we adopt the popular Proximal Policy Optimization (PPO) algorithm from reinforcement learning to learn an allocation policy by training a policy and value function network. We demonstrate in simulations that our approach outperforms the best heuristic baseline by up to a factor of 4.

## 1 Introduction

Each year, more than 250,000 women lose their lives during and following pregnancy and childbirth (World Health Organization 2024), with the first 24 hours post-delivery being particularly perilous (Li et al. 1996; Dol et al. 2022). A significant contributing factor to this tragic statistic is the poor quality of care available in under-resourced communities (Crear-Perry et al. 2021). Consequently, there is growing interest in different ways of improving peripartum care to prevent life-threatening complications such as hemorrhage, hypertensive disorders, and sepsis. One key approach is through the monitoring of maternal vital signs, which can be used to identify complications early on via early warning systems that provide an opportunity for timely clinical intervention (Vousden et al. 2019).

In fact, the World Health Organization (2016) recommends close monitoring of maternal vital signs in the first 24



Figure 1: Wireless vital sign monitoring device on a mother.

hours after birth, thereby highlighting the importance of maternal vital signs in high-quality maternal care. Traditionally, this monitoring is conducted by healthcare providers who manually measure vital signs at regular intervals. However, even in well-resourced settings, the recommended monitoring frequency poses a substantial burden. In resource-limited settings, meeting these guidelines has been very difficult to achieve (Mugenyi et al. 2021; Semrau et al. 2017).

An automated alternative to measuring vital signs is the use of wearable or wireless vital sign monitoring devices, often in the form of wireless biosensors; see Figure 1 for a picture of such a device (Boatin et al. 2016, 2023). These sensors continuously measure and transmit the mother's vital signs, with the option to trigger alerts if abnormalities in vital signs are detected. These alerts give clinicians the opportunity to initiate appropriate medical responses in real time when needed. Thus, automated vital sign monitoring using wireless biosensors provides an opportunity to implement early warning systems in resource-constrained environments, where to date, human resource constraints have limited the ability to check vital signs consistently. Such systems have been demonstrated to be functional and acceptable in these settings (Boatin et al. 2016; Ngonzi et al. 2017). However, in practice, while these monitoring devices enhance human resources, their availability will still be severely limited, leading to the central research question addressed in this paper: *Who should wear a monitoring device and for how long?*

While our research is motivated by several hospital settings (in Uganda and Ghana), the specific motivation for this study comes from the Department of Obstetrics and

Gynecology at the Mbarara Regional Referral Hospital in Mbarara, Uganda. With approximately 9000 deliveries annually, this hospital serves as the primary referral center for southwestern Uganda. It has been piloting the use of wireless vital sign monitors during the critical 24 hour period after birth (Boatin et al. 2021, 2023; Mugenyi et al. 2021).

We tackle the problem of allocating monitoring devices by framing it as an instance of the popular restless multi-armed bandit (RMAB) model. This model has been successfully applied to distribute limited resources in different healthcare contexts (Ayer et al. 2019) as well as in areas beyond healthcare, such as anti-poaching (Qian et al. 2016) and machine maintenance (Abbou and Makis 2019). However, unique challenges arise in our application, rendering existing solution methods inadequate. Specifically, in contrast to classic RMAB problems, our allocation setting places several novel constraints on the allocation, for instance, each mother must be allocated a device for a minimum and maximum duration, and once a device is removed from a mother, it cannot be (easily) reassigned to her at a later point (see Section 2). Furthermore, currently, no historical data or other relevant features of arriving mothers are available to the algorithm in the intended deployment setting. Consequently, the decision of when to remove a device after the minimum monitoring period must be based solely on the mother’s so-far recorded vital signs. Therefore, our algorithm learns in an online manner, leveraging information collected from previously monitored mothers to make decisions for new ones. To address these challenges, we adopt the Proximal Policy Optimization (PPO) algorithm, a reinforcement learning technique that has proven effective in diverse domains (Schulman et al. 2017) such as video games (Kristensen and Burelli 2020), robotics (Melo, Melo, and Maximo 2021), and autonomous vehicles (Guan et al. 2020). Our approach leverages PPO’s strengths in learning robust policies under complex constraints and dynamic environments. In sum, our main contributions are<sup>1</sup>:

- We are the first to identify and formalize the algorithmic problem of allocating scarce wireless vital sign monitoring devices, with novel real-world constraints previously unstudied by the resource allocation literature.
- We develop a modern reinforcement learning-based algorithm to learn and make allocation decisions in real time using RMABs, contributing to both the application domain and extending the RMAB literature.
- We demonstrate in simulation that our algorithm significantly outperforms natural heuristic baselines, achieving improvements ranging from 100% to 400%, which gives a good indication of the usefulness of AI approaches for this problem. Moreover, we conduct a first analysis on limited data from the Mbarara Regional Referral Hospital and describe the next steps toward responsible real-world AI deployment, including additional data collection, quality control, ethics reviews, and field trials.

---

<sup>1</sup>Additional details on our experimental setup and results can be found in our full version (Boehmer et al. 2024).

**Related Work on RMABs** There is a rich body of work on devising allocation algorithms for scarce resources using the restless bandits model. Different variants of this model have been studied, differing in the information available to the planner (Chen and Hou 2024), the constraints on budget usage (Rodriguez-Diaz et al. 2023), and resource allocation strategies (Mao and Perrault 2024). The work of Mate et al. (2022) studies streaming restless multi-armed bandits, where arms appear and disappear over time, which is also the case in our problem. The primary distinction between their setting and ours lies in their assumption that arms have discrete states with known transition dynamics, whereas in our problem, transition dynamics are unknown and states are not restricted to being discrete. Moreover, we have additional constraints placed on the allocation.

In contrast, Zhao et al. (2024) propose a RL-based solution for streaming bandits where transition dynamics are unknown and states consist of a single continuous value. However, their method is not applicable to our setting, as it cannot accommodate the additional allocation constraints specific to our problem (specifically, the lambda network they use is incompatible with constraints beyond the standard budget one). Moreover, their approach relies on feature information that is unavailable in our context and is limited to simpler state spaces, which do not adequately capture the complexity of vital sign profiles (see Section 2).

## 2 Application, Modeling, and Challenges

### 2.1 Application Details

We outline the characteristics of the problem encountered in our application domain. Our setting is the maternity unit of the Mbarara Regional Referral Hospital (Boatin et al. 2021, 2023): Mothers arrive at the maternity unit and deliver throughout the day. After delivery, mothers remain in the maternity unit for some time before being redirected to other care measures or discharged. Mothers will wear a monitoring device during some time they spend in the unit.

Wearing a monitoring device has no direct impact on the mother’s vital signs or health. However, there is a clear indirect impact: If a monitored vital sign deviates from a pre-set normal range, an alert is sent to the responsible clinician’s phone. While some alerts may be disregarded due to capacity constraints or other factors, in most cases, the clinician will visit the patient, manually assess any abnormal vital signs, and initiate appropriate clinical interventions if needed. These interventions are expected to positively influence the mother’s health and stabilize her vital signs.

There are several external constraints imposed on the allocation of monitoring devices. First, every mother should wear the device for a minimum duration, for instance, the initial two hours after birth, which are particularly high-risk. Second, once a device is removed from a mother, it cannot be (easily) reassigned to her, as she will transition to a different set of care protocols. Third, each mother is only eligible to wear the device for the first 24 hours after birth, as this is the targeted monitoring period for the program. No feature information about the mother—such as demographic details or historical medical data—is available to the algorithm, con-

stituting a safeguard for data privacy.<sup>2</sup> Consequently, the decision when to remove the device from a mother after her minimum monitoring period has ended must be made purely based on her previously recorded vital signs.

The objective of the allocation strategy is to minimize occurring complications and maintain the vital signs of all patients—whether they are wearing a device or not—within the normal range during their stay in the maternity unit. This implies that patients at higher risk of developing abnormal vital signs indicating potential complications should be prioritized: For them, a monitoring device will trigger alerts and prompt the needed timely clinical assistance.

## 2.2 Formal Modeling & Challenges

We model the problem of allocating monitoring devices using the RMAB framework. An instance consists of a planning horizon  $T$ , a budget  $B$  (the number of available devices), a discount factor  $\gamma$ , a minimum  $t_{\min}$  and maximum  $t_{\max}$  number of steps a mother should be monitored, and a set  $N$  of mothers (called *arms*). Each arm  $i \in N$  follows a Markov Decision Process  $(\mathcal{S}_i, \mathcal{A}_i = \{0, 1\}, \Gamma_i, R_i)$ .  $\mathcal{S}_i$  represents the possible *states* of arm  $i$ . In our application, states are multidimensional and continuous and include the current values of the vital signs, along with potentially aggregated statistics like the variability of each vital sign over recent time steps.  $\mathcal{A}_i = \{0, 1\}$  represents the actions, where 0 denotes the *passive* action and 1 denotes the *active* action, i.e., allocate a device to the mother.  $\Gamma_i$  describes the parameters characterizing how arm  $i$ 's state evolves from one step to the next conditioned on the taken action.<sup>3</sup>  $R_i : \mathcal{S}_i \rightarrow \mathbb{R}$  is the reward function of arm  $i$ , penalizing states where vital signs fall significantly outside the normal range, indicating potential complications. Additionally, each mother has an arrival  $\alpha_i \in [T]$  and departure time  $\beta_i \in [T]$ . We assume that the state space  $\mathcal{S}$  and reward function  $R$  are the same for each arm and are known. In contrast,  $\Gamma_i$  is arm-specific and unknown. Additionally,  $\alpha_i$  and  $\beta_i$  are also arm-specific and are revealed at the respective timestep. At each timestep  $t \in [T]$ , an arm  $i$  is *present* if  $t \in [\alpha_i, \beta_i]$ . Let  $N_t$  be the set of arms present at time  $t$ . The goal is to learn a policy  $\pi$  that maps the set of currently present arms and their current states  $\mathbf{s} \in \mathcal{S}^{|N_t|}$  to an action vector  $\mathbf{a} \in \{0, 1\}^{|N_t|}$ , such that for each  $t \in [T]$  and  $i \in N_t$ : (1)  $\sum_{j \in N_t} a_j \leq B$ , (2)  $a_i = 1$  if  $t \in [\alpha_i, \alpha_i + t_{\min} - 1]$ , (3)  $a_i = 0$  if  $t > \beta_i$  or  $t \geq \alpha_i + t_{\max}$ , and (4)  $a_i = 0$  if there is some step  $t' \in [\alpha_i, t]$  in which  $i$  was assigned the passive action.

<sup>2</sup>We note that in some hospitals such information exists on paper and could be digitalized if needed. However, an algorithm that does not require features and historical data is naturally much easier to deploy and preferable from a privacy and safety perspective.

<sup>3</sup>For example, transitions of continuous states might follow multivariate Gaussian distributions (see Section 4.1), with separate distributions for the active and passive action. Then,  $\Gamma_i$  contains the distributions' mean and covariance matrix. In many states, applying the active action (i.e., allocating a device) will not alter the transition dynamics unless an alert is triggered. Nonetheless, it is beneficial to continue monitoring such arms as they may transition into critical states later on where alerts are generated and the transition dynamics are impacted by the active action.

The goal is to find such a policy that maximizes the accumulated discounted reward:  $\sum_{t \in [T]} \gamma^{t-1} \mathbb{E}_{\mathbf{s} \sim (N, \pi)} \sum_{i \in N_t} R(s_i)$ . Note that due to the allocation constraints, allocation decisions only need to be made when a new arm arrives. At that moment, the algorithm must assign the active action to the new arm (due to the minimum monitoring period). The “only” decision the algorithm needs to make is which arm should be flipped from the active to the passive action, i.e., from which mother we take the monitoring device needed for the new mother.

**Novel Challenges** Our problem introduces three novel aspects that set it apart from existing work in restless bandits:

- The standard assumption in the restless bandit literature is that states are few and discrete, which simplifies transition dynamics (Niño-Mora 2023). However, vital signs evolve in complex, continuous ways, prohibiting the discretization of the state space. Further, patient's states are characterized by multiple continuous vital sign values plus statistics about their trajectory.
- Existing works on RMABs with unknown transition probabilities typically rely on arm's feature information to learn their dynamics. In our context, no such features are available, forcing the algorithm to estimate an arm's future behavior based solely on its current state.
- Allocation constraints (2)-(4) from above are important for many monitoring applications but have seemingly not been previously explored in the RMAB literature.

## 3 Methodology

To address the novel challenges arising in our application domain, we employ an actor-critic approach using Proximal Policy Optimization (PPO) for policy updates, which is widely used in reinforcement learning. Our algorithm requires access to a simulator “Simulate( $i, s_i, a_i$ )” of the environment that takes as input an arm, its current state, and its action and outputs the new state of the arm. The idea is to train a policy using Algorithm 1 which has access to the simulator and then deploy the learned policy in the real-world, which ensures the required high-quality decision-making from the start (Zhao et al. 2024).

The actor is a policy neural network that takes as input the current state of a single arm and outputs the action probability for both possible actions for this arm. We let  $\theta(a | s)$  denote the action probability for action  $a \in \{0, 1\}$  returned by the network on input  $s \in \mathcal{S}$ . Our policy acts on the arms for which the policy network returns the highest probability for the active action. Thus, the outputs of the policy network determines which arms are assigned the active action. The critic is another neural network that takes as input the state  $s \in \mathcal{S}$  of an arm and outputs the baseline estimate  $V(s)$ , which is the expected total discounted reward generated by this arm starting from state  $s$  assuming that actions are taken following the action probabilities returned by the policy network. The critic network is used for updating the policy network via the PPO algorithm. Importantly, both networks are shared among all arms, enabling arms to learn from each

---

**Algorithm 1: RL-based allocation algorithm**

---

```
1: Input:  $n_{\text{epoch}}$  instances of our problem.
2: Initialize actor  $\theta$  and critic  $\phi$ 
3: for each of the  $n_{\text{epoch}}$  instances do
4:   for  $t = 1, \dots, T$  do
5:     Let  $N_t^{\text{new}} := \{i \in N_t \mid \alpha_i + t_{\text{min}} > t\}$ 
6:     Assign the active action to all arms from  $N_t^{\text{new}}$ 
7:     Let  $N_t^{\text{eligible}}$  be the set of all arms  $i \in N_t$  with
       $\alpha_i + t_{\text{max}} < t$  and which have never been assigned the
      passive action since they arrived in step  $\alpha_i$ 
8:     Use policy network to compute action probability
       $\theta(a_i \mid s_i)$  for each arm  $i \in N_t$ 
9:     Assign the active action to  $B - |N_t^{\text{new}}|$  arms
      from  $N_t^{\text{eligible}} \setminus N_t^{\text{new}}$  with the highest active action
      probability  $\theta(1 \mid \cdot)$  and the passive action to all others
10:    for  $i \in N_t$  do
11:       $s'_i = \text{Simulate}(i, s_i, a_i)$ 
12:      Add tuple  $(s_i, a_i, R(s_i), s'_i)$  to buffer
13:       $s_i \leftarrow s'_i$ 
14:    Update actor-critic  $(\theta, \phi)$  pair via PPO using buffer
```

---

other—this is crucial because each arm remains in the system for only a limited time.

Algorithm 1 proceeds in multiple epochs. For training purposes, each epoch deals with a separate set of arms. Within each epoch, a fixed policy is used to make the allocation decisions while respecting all constraints. The behaviors of all arms are observed and stored in a buffer. At the end of each epoch, we use the buffer to update the policy and critic networks, thereby refining the policy.

Breaking down Algorithm 1, in Lines 5–9, the algorithm assigns actions to all present arms: In Line 6, the algorithm assigns the active action to all arms that have been in the system for less than  $t_{\text{min}}$  steps, as they have not been monitored for the required minimum time. In Line 9, the remaining active actions are assigned to the arms eligible for receiving an active action in this step with the highest action probability returned by the policy network. Then, in Line 11, we simulate the next state of each arm conditioned on the assigned action and update its state accordingly in Line 13.

The policy and critic networks are updated following the principles of the PPO algorithm (Schulman et al. 2017). Specifically, let  $V(s)^j$  be the values returned by the critic network at the end of epoch  $j$ . We compute the advantage function  $A^j(s, a)$  for epoch  $j$ , for  $s \in \mathcal{S}$  and  $a \in \{0, 1\}$ , which quantifies the benefit of taking a certain action  $a$  in state  $s$  compared to the current policy as  $A^j(s, a) = Q^j(s, a) - V^j(s)$ , where  $Q^j(s, a)$  is the expected discounted cumulative reward for the completion of the current episode under the current policy for an arm in state  $s$  to which action  $a$  is applied in this step. The advantage function is then incorporated into the actor’s policy gradient to update the policy network in the actor, following the standard PPO procedure. When running the algorithm in testing, we execute Lines 4 to 13 with the trained policy network.

## 4 Experiments

We present our experiments using data from a publicly accessible vital sign database, which offers rich and high-quality data. Our goal is to demonstrate the general capabilities of our algorithm; we revisit our initial use case of maternal care in Section 5. All experiments are conducted in simulation and only work with simulated vital sign profiles. In Section 4.1, we describe our setup, including the datasets used, the instance generation, the trained simulator, and the baselines employed. In Section 4.2, we present the results.

### 4.1 Setup

**Domain** The experiments in this section are based on data from the widely used MIMIC-III (Johnson et al. 2016) and MIMIC-IV (Johnson et al. 2023) datasets, which have become popular and influential in computer science research. Both datasets contain de-identified clinical data from thousands of patients who stayed in critical care units at Beth Israel Deaconess Medical Center in Boston over different periods, including vital sign measurements typically recorded at one-hour intervals. In our experiments, each arriving arm corresponds to a new patient entering the critical care unit. For MIMIC-III, we use the vital signs (i) heart rate, (ii) speed of breathing (respiratory rate), and (iii) blood oxygen saturation (SPO2), while for MIMIC-IV, we use (i) heart rate, (ii) respiratory rate, and (iii) skin temperature. We normalize each vital sign between 0 and 1 using min-max normalization. For each patient, we take the median vital sign value at each hour. Thus, one timestep corresponds to one hour. We exclude patients with fewer than 10 data points.

**Simulator** Patients’ state representation includes, for each vital sign, its current value and its variance over the last five timesteps. The reward function assigns a reward of 0 if all vital sign values fall within a normal range (Boatin et al. 2021). For each abnormal vital sign, we incur a negative reward that shrinks exponentially with the extent of the deviation from the normal (see (Boehmer et al. 2024) for details). The exponential penalties model the increasing severity associated with larger deviations from the normal range.

The patient’s behavior is governed by a multivariate Gaussian distribution defined over the vital sign values in the current step and in the next step. We sample the initial state of a patient from this distribution by taking a sample and using only the sampled vital sign values in the current step. Under the passive action, the next state is sampled from the conditional distribution of the Gaussian, given the current state. Under the active action, we make a case distinction. If all vital signs are within the normal range, the state transitions as under the passive action since no alert is triggered, and the device does not influence the patient’s trajectory. If any vital sign is abnormal, with a probability of 30%, the state transitions as under the passive action (modeling cases where clinicians do not respond to the alert, which occur approximately 30% of the time (Boatin et al. 2021, 2023)). Otherwise, the abnormal vital signs are probabilistically adjusted towards the normal range before sampling the next state, reflecting the positive effects of clinical intervention following an alert (see (Boehmer et al. 2024) for details).

## 4.2 Results

**Instances** We set  $T = 100$ ,  $t_{\min} = 3$ ,  $t_{\max} = 25$ , and assume that patients leave 50 steps after they join. We vary the budget  $B$ . Initially, there are  $B$  patients, and every five steps new patients join. We report the number of patients  $N$ , which describes the “typical” number of patients in the system. The inner workings of our instances are best understood by looking at a concrete example: Let us consider  $B = 3$  and  $N = 20$ . Initially, there are  $B = 3$  patients, and every five timesteps, two (i.e.,  $N/10$ ) new patients arrive. Since each patient leaves 50 steps after they join and every five steps two patients arrive, the number of patients in the system gradually grows to 20 and stabilizes at 20. Importantly, at every point at which the algorithm makes an allocation decision, there are only few patients to pick from. In this instance, there are five: two (i.e.,  $N/10$ ) newly arriving patients and three (i.e.,  $B$ ) existing patients who currently hold a device.

It remains to describe how we sample patients’ transition parameters. For this, we fit a weighted mixture of five multivariate Gaussians (i.e. five components) on the dataset, where we partition all trajectories into tuples that include the vital signs in the current and next step. When we sample a patient, we first select a component from the Gaussian mixture based on the component’s weights. This determines the initial mean and covariance of the patient. To introduce variability, we linearly combine this mean and covariance with those of another randomly selected component, using a weight uniformly sampled between 0 and 0.15.

**Baselines** We refer to NoAction as the policy that does not allocate any monitoring devices. All other baselines respect the allocation constraints. Recall that this means that they only need to make a decision if a new patient appears. Then, the algorithm needs to decide from which patient currently holding a device and having already been monitored for  $t_{\min}$  steps, we remove the device.

- **Random:** Randomly selects a patient.
- **Extreme Values:** Selects the patient with the least abnormal vital signs, i.e., the patient with the most extreme summed normalized vital sign values.
- **Highest Variability:** Selects the least stable patient, i.e., the patient with the highest summed variance of vital signs over the past five steps.

Extreme Values is particularly intuitive, as it takes the devices away from the patients generating the highest reward.

**Training and evaluation** For each considered setting, we average our results over 100 seeds. For each seed, we do the following: We train our algorithm for  $n_{\text{epoch}} = 50$  epochs. At the beginning of each epoch, we create a new instance by sampling a fresh set of  $N$  arms. Subsequently, we evaluate the trained policy along with various baselines on 50 newly generated instances and compute the average rewards. Both in testing and evaluation, patients transition according to the simulator as described in Section 4.1.

**Results** See Figure 2 for an overview of our results, where we vary the budget  $B$  and the number of patients  $N$ . We report the reward averaged over 100 random seeds, where

we normalize the reward of the methods by subtracting the reward of the No Action baseline and dividing by  $N$ . Consequently, the reported values capture the benefit of the allocated monitoring devices. Our method outperforms the baselines across all examined settings. Notably, when  $B = 3$  and  $N = 20$ , we outperform the best baseline by 433% and 173% on the MIMIC-II and MIMIC-IV datasets, respectively; when  $B = 6$  and  $N = 50$ , we outperform the best baseline by 431% and 141% on the MIMIC-III and MIMIC-IV data, respectively. Interestingly, the intuitive Extreme Values and Highest Variability baselines perform worse than Random. The fact that both baselines are insufficient highlights the complexity of our problem and motivates the necessity for a more intricate approach like ours. We also observe that devices are typically taken away when a patient’s vital signs are within a medium range and have low variability (Boehmer et al. 2024).

## 5 Maternal Care in Mbarara: Initial Results

We revisit the specific application of maternal care in Mbarara discussed in Section 2.1. We have access to continuous vital sign measurements from 100 patients collected using monitoring devices at the Mbarara Regional Referral Hospital.<sup>4</sup> As in Section 4, we discretize the trajectories into 60-minute intervals by taking the median value of each vital sign recorded during the interval. The vital signs we consider are: (i) heart rate, (ii) respiratory rate, (iii) and blood oxygen saturation (SPO2). The rest of the setup is as in Section 4.

Note that we can only fit the simulator on the 100 patients’ traces available to us (we are working closely together with Mbarara University of Science and Technology to collect more data and build a more powerful simulator). As a result, the patients’ behaviors captured by the simulator are quite simple and less stable compared to the much larger MIMIC datasets. In the absence of complex patterns and interactions in the vital signs generated by the simulator, simpler algorithms (e.g., the random strategy) are expected to show improved behavior. *A key question we wish to address in this section is whether the limited dataset already provides evidence that a purely random strategy is insufficient and a more involved approach is needed?*

Figure 3 shows the results in three different settings. We observe that our method significantly outperforms the Random as well as the Highest Variability strategies. Furthermore, while the difference with the second-best Extreme Values is not statistically significant, there is still

---

<sup>4</sup>These trajectories were collected as part of a hybrid effectiveness-implementation trial of a wearable vital sign monitor among post-cesarean women. Women in this trial consented to wear the monitor for 24 hours. Ethics approval for this trial was obtained from the Mbarara University of Science and Technology Research Ethics Committee (17/10-18), the Uganda Council of Science and Technology (HS417ES) and the MassGeneralBrigham Institutional Regulatory Board (2019P000095). The trial was registered at clinicaltrials.gov (NCT04060667). We use the previously collected vital sign trajectories from trial participants to develop a simulator, however, vital sign data generated in these simulations do not represent and cannot be traced to real participants.




Figure 2: Results on MIMIC-III (top) and MIMIC-IV (bottom), averaged over 100 random seeds for varying budget  $B$  and number of patients  $N$ . The error bars show the standard error of the generated reward, which is normalized by subtracting the reward of the No Action baseline and then dividing by  $N$ .




Figure 3: Results on data from the Mbarara Hospital, averaged over 100 random seeds. Error bars show the standard error of rewards, which are normalized by subtracting the reward of the No Action baseline and then dividing by  $N$ .

a trend of our algorithm showing superiority, as it outperforms this best baseline by 29% and 20% for settings ( $B = 3, N = 20$ ) and ( $B = 7, N = 40$ ), respectively. In light of the results from Section 4.2, we expect these differences to grow significantly if we add additional training data for the simulator and in real-world deployment. Thus, our results provide a first evidence for the advantages of our method in the maternal care domain.

## 6 Path to Deployment

While our experiments demonstrate the potential of using RL-based algorithms for allocating monitoring devices, several important steps remain before real-world deployment. We are currently planning to collect additional vital sign trajectories in the Mbarara Regional Referral Hospital based on which we want to refine our model, especially regarding the impact of wearing a device. Once the simulator is trained on a larger and more diverse dataset, we will conduct a rigorous evaluation of the simulator and the learned policy, including assessing potential biases, verifying policy convergence and its robustness to distribution shifts in patient populations, and making necessary adjustments. Once the policy’s decision-making is thoroughly validated, we will proceed with obtaining ethics and regulatory approval to test the policy in a real-world setting. Recall that, as a safeguard for data privacy, no feature information about the mother is available to the algorithm. Next, we will run a first, continuously closely monitored trial in the Mbarara Regional Referral Hospital to test the implementation pipeline, safety, and

acceptability of our method, and to conduct a preliminary analysis of its effectiveness. Assuming this study meets pre-defined milestones regarding feasibility, acceptability, and safety, the final phase consists of a comprehensive evaluation through a randomized controlled trial (RCT) in multiple hospitals. At the conclusion of a successful RCT, we will focus on the careful, responsible deployment of our system. Throughout this entire path to deployment, we will maintain a very close collaboration with domain experts and agencies to thoroughly check for biases, and ensure steps towards a responsible deployment. Additionally, we plan to explore the broader application of RL-based algorithms in other post-surgical care settings where monitoring devices can be used to improve patient outcomes.

## 7 Conclusion

We identified the problem of distributing wireless vital sign monitoring devices—particularly relevant in peripartum maternal care—as a novel resource allocation challenge. We introduced an RMAB-style model for this problem, which differs from previously studied models in several key aspects. Our experiments demonstrate that our RL-based allocation algorithm enables more efficient use of limited monitoring devices. There are several promising directions for future research. The first is the path outlined in Section 6. Beyond this, the unique characteristics of our setting motivate the study of new variants of RMAB models. For instance, it would be interesting to develop algorithms with performance guarantees for traditional RMAB settings where arms’ MDPs are discrete and known but some of our allocation constraints must be respected. Our application also raises additional algorithmic questions. One notable challenge is optimizing the design of alerts sent by monitoring devices: While sending more alerts increases the likelihood of an alert being sent before or during a complication, it also increases the burden on clinicians and reduces their responsiveness to each individual alert. An intricate challenge for future work is determining optimal thresholds for vital sign alerts that strike the right balance between timely detection of complications and minimizing alert fatigue.

## Acknowledgments

This work is supported by Harvard Data Sciences Institute and by the AI Research Institutes Program funded by the National Science Foundation under AI Institute for Societal Decision Making (AI-SDM), Award No. 2229881. Research was sponsored by the DARPA and was accomplished under Grant Number: HR001122C0182.

## References

- Abbou, A.; and Makis, V. 2019. Group maintenance: A restless bandits approach. *INFORMS J. Comput.*, 31(4): 719–731.
- Ayer, T.; Zhang, C.; Bonifonte, A.; Spaulding, A. C.; and Chhatwal, J. 2019. Prioritizing hepatitis C treatment in US prisons. *Oper. Res.*, 67(3): 853–873.
- Boatin, A. A.; Del Cueto, P.; Bebell, L. M.; et al. 2023. Implementation of wireless continuous vital sign monitoring after cesarean delivery in Uganda. *Am. J. Obstet. Gynecol.*, 228(1): S246.
- Boatin, A. A.; Ngonzi, J.; Wylie, B. J.; et al. 2021. Wireless versus routine physiologic monitoring after cesarean delivery to reduce maternal morbidity and mortality in a resource-limited setting: protocol of type 2 hybrid effectiveness-implementation study. *BMC Pregnancy and Childbirth*, 21: 1–12.
- Boatin, A. A.; Wylie, B. J.; Goldfarb, I.; Azevedo, R.; Pittel, E.; Ng, C.; and Haberer, J. E. 2016. Wireless vital sign monitoring in pregnant women: a functionality and acceptability study. *Telemed. e-Health*, 22(7): 564–571.
- Boehmer, N.; Zhao, Y.; Xiong, G.; Diaz, P. R.; Cibrian, P. D. C.; Ngonzi, J.; Boatin, A.; and Tambe, M. 2024. Optimizing Vital Sign Monitoring in Resource-Constrained Maternal Care: An RL-Based Restless Bandit Approach. [abs/2410.08377](https://arxiv.org/abs/2410.08377).
- Chen, X.; and Hou, I.-h. 2024. Contextual Restless Multi-Armed Bandits with Application to Demand Response Decision-Making. In *Proc. of CDC 2024*.
- Crear-Perry, J.; de Araujo, R. C.; Johnson, T. L.; McLemore, M. R.; Neilson, E.; and Wallace, M. 2021. Social and Structural Determinants of Health Inequities in Maternal Health. *J. Women's Health*, 30(2): 230–235.
- Dol, J.; Hughes, B.; Bonet, M.; et al. 2022. Timing of maternal mortality and severe morbidity during the postpartum period: a systematic review. *JBI evid. synth.*, 20(9): 2119–2194.
- Guan, Y.; Ren, Y.; Li, S. E.; Sun, Q.; Luo, L.; and Li, K. 2020. Centralized Cooperation for Connected and Automated Vehicles at Intersections by Proximal Policy Optimization. *IEEE Trans. Veh. Technol.*, 69(11): 12597–12608.
- Johnson, A. E.; Bulgarelli, L.; Shen, L.; et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Sci. data*, 10(1): 1.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Sci. data*, 3(1): 1–9.
- Kristensen, J. T.; and Burelli, P. 2020. Strategies for Using Proximal Policy Optimization in Mobile Puzzle Games. In *Proc. FDG 2020*.
- Li, X. F.; Fortney, J. A.; Kotelchuck, M.; and Glover, L. H. 1996. The postpartum period: the key to maternal mortality. *Int. J. Gynecol. Obstet.*, 54(1): 1–10.
- Mao, Y.; and Perrault, A. 2024. Time-Constrained Restless Multi-Armed Bandits with Applications to City Service Scheduling. In *Proc. of AAMAS 2024*, 2375–2377.
- Mate, A. S.; Biswas, A.; Siebenbrunner, C.; Ghosh, S.; and Tambe, M. 2022. Efficient Algorithms for Finite Horizon and Streaming Restless Multi-Armed Bandit Problems. In *Proc. of AAMAS 2022*, 880–888.
- Melo, L. C.; Melo, D. C.; and Maximo, M. R. O. A. 2021. Learning Humanoid Robot Running Motions with Symmetry Incentive through Proximal Policy Optimization. *J. Intell. Robot. Syst.*, 102(3): 54.
- Mugenyeni, G. R.; Ngonzi, J.; Wylie, B. J.; Haberer, J. E.; and Boatin, A. A. 2021. Quality of vital sign monitoring during obstetric hospitalizations at a regional referral and teaching hospital in Uganda: an opportunity for improvement. *Pan Afr. Med. J.*, 38(1).
- Ngonzi, J.; Boatin, A.; Mugenyeni, G. R.; Wylie, B. J.; and Haberer, J. E. 2017. A functionality and acceptability study of wireless maternal vital sign monitor in a tertiary university teaching hospital in rural Uganda. *J. Women's Health Gynecol.*, 1: 1–8.
- Niño-Mora, J. 2023. Markovian Restless Bandits and Index Policies: A Review. *Mathematics*, 11(7).
- Qian, Y.; Zhang, C.; Krishnamachari, B.; and Tambe, M. 2016. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In *Proc. of AAMAS 2016*, 123–131.
- Rodriguez-Diaz, P.; Killian, J. A.; Xu, L.; Suggala, A. S.; Taneja, A.; and Tambe, M. 2023. Flexible Budgets in Restless Bandits: A Primal-Dual Algorithm for Efficient Budget Allocation. In *Proc. of AAAI 2023*, 12103–12111.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Semrau, K. E.; Hirschhorn, L. R.; Marx Delaney, M.; et al. 2017. Outcomes of a coaching-based WHO safe childbirth checklist program in India. *N. Engl. J. Med.*, 377(24): 2313–2324.
- Vousden, N.; Lawley, E.; Nathan, H. L.; Seed, P. T.; et al. 2019. Effect of a novel vital sign device on maternal mortality and morbidity in low-resource settings: a pragmatic, stepped-wedge, cluster-randomised controlled trial. *The Lancet Global Health*, 7(3): e347–e356.
- World Health Organization. 2016. Standards for improving quality of maternal and newborn care in health facilities.
- World Health Organization. 2024. Maternal Mortality. Accessed: 2024-08-28.
- Zhao, Y.; Behari, N.; Hughes, E.; Zhang, E.; Nagaraj, D.; Tuyls, K.; Taneja, A.; and Tambe, M. 2024. Towards a Pre-trained Model for Restless Bandits via Multi-arm Generalization. In *Proc. of IJCAI 2024*, 321–329.