AI for Planning Public Health Interventions

Aditya Mate
Harvard University
aditya_mate@g.harvard.edu

Abstract

Several scenarios involving public health interventions, have an unifying underlying theme, that deals with the challenge of optimizing the limited intervention resources available. My dissertation casts this as a Restless Multi-Armed Bandit (RMAB) planning problem, identifying and addressing several new, fundamental questions in RMABs.

1 Introduction

Community Health Workers (CHWs) are members of the local community, who serve as the link between primary healthcare centers and the people of the local community. They deliver a wide range of public health services, such as monitoring health, screening, basic emergency care, health education, delivering interventions as necessary, etc. However, a key constraint that these health workers grapple with is that their intervention resources may be severely limited. For example, health workers in India monitoring tuberculosis patients, operate by calling patients in their cohort each day to monitor their adherence to prescribed medication and to simultaneously deliver interventions over the phone to encourage adherence among patients. Each such health worker may be routinely responsible for managing upto 200 patients or even more, making it infeasible to call and intervene on each patient on a daily basis. In such a scenario, the health workers must prioritize patients and plan the interventions smartly.

Towards addressing this intervention planning problem, one key contribution of my research is to cast it as a Restless Multi-Armed Bandit (RMAB) problem. RMAB is a popular framework that has been used to address limited resource planning problems in the past, and has been employed in a myriad of domains such as sensor maintenance, anti-poaching patrol planning, multi-channel communication [Liu and Zhao, 2010], age of information [Sombabu *et al.*, 2020], etc. However, because no known solution to general RMAB problems is available, previous studies have focused on specific problem instances to derive results restricted to that specific problem. Most previous work in RMAB either fails to capture all the complexities involved in the health monitoring and intervention planning problem faced by the health workers or are computationally too expensive and thus very inconvenient to use. This opens up new fundamental challenges in the space of restless bandits and calls for innovative solutions.

Figure 1: Community Health Worker (CHW) engaging with patients and delivering interventions. Picture Credits: Pippa ranger

My dissertation introduces a new subclass of RMABs that models this intervention planning problem. My research harnesses the special structure of the problem to design fast algorithms that scale well, without having to sacrifice on performance [Mate *et al.*, 2020]. My research also considers several other challenges that health workers encounter which planning such interventions and has proposed solutions to account for those [Mate *et al.*, 2021a; Mate *et al.*, 2021b]. Section 2 formalizes the intervention planning problem setup and section 3 outlines the key contributions of my research towards solving these challenges.

2 Problem Setup

We cast the health monitoring and intervention planning problem as a restless multi-armed bandit (RMAB) consisting of a set of N arms (denoting patients) that the planner must monitor. Each arm follows a Markov Decision Process (MDP) that can be in one of two states – a ‘good’ state and a ‘bad’ state. The process keeps transitioning between the states each time step, according to a fixed, known transition matrix. This process goes on for a horizon length of T time steps. Further, the planner has access to k intervention resources per time step, where $k \ll N$. The transition dynamics are such that the intervention action has a higher propensity to push the arm to the ‘good’ state. The planner must thus decide, which k out of the N arms to pull each time step, to maximize the benefit of the limited intervention resources.

3 Overview of Contributions

My research adopts the Whittle Index solution technique, proposed by [Whittle, 1988], that has been the dominant paradigm used for solving RMAB problems, and has been shown to perform very well. However, success of the Whittle Index approach hinges on two key challenges that must be overcome: (1) The approach has been shown to be asymptotically optimal only if a technical condition called 'indexability' is satisfied. There is no general technique available for verifying indexability for any arbitrary RMAB, and hence is usually difficult to show. (2) Secondly, computation of the index itself is challenging and can be computationally intensive. Previous approaches have proposed techniques towards speeding up the index policy computation, but the method still remains computationally intensive.

3.1 Collapsing Bandits

My research proposes *Collapsing Bandits* [Mate *et al.*, 2020], a subclass of RMAB and leverages its structure to prove key theoretical results that (1) prove indexability (2) enable us to compute the Whittle Index with a dramatic speed-up. Specifically, we prove that a specific subclass of RMAB policies, called 'Threshold Policies' is optimal under easily verifiable conditions, which in turn can be shown to imply indexability. Further, optimality of 'Threshold Policies' also allows efficient computation of the Whittle Index, achieving a 3-order-of-magnitude speed-up for the motivating task of tuberculosis medication adherence, evaluated using a real-world data set. Moreover, the algorithm achieves this speedup without having to sacrifice on performance.

3.2 Risk-Aware Bandits

The planning algorithm proposed, however, optimizes for the average adherence of the patient cohort and may be inadequate to handle a real-world planner's complex considerations. For instance, the algorithm may be insensitive to equitable resource allocation – it may sometimes deem some patients to be less valuable to intervene on and may ignore them completely. Similarly it relies on receiving perfect observations of the patients' adherence upon intervening and assumes the planner to be risk-neutral, even though real-world health workers may be risk-averse. To address these concerns, my research proposes a Risk-Aware algorithm [Mate *et al.*, 2021b] that can admit more general, non-linear reward functions that can be suitably shaped to accommodate the planner's complex considerations while also taking imperfect observations into account. Essentially, the new algorithm is able to optimize for utility measured according to a yardstick of the planner's choice rather than optimizing for aggregate statistics, such as the average adherence of the patient cohort.

3.3 Streaming Bandits

The Collapsing Bandits framework proposes an algorithm that works well; however it makes a couple of limiting assumptions. It proposes an algorithm designed for infinite horizon settings and relies on the duration of the health program being sufficiently large. However, health programs may be of finite length in the real-world and sometimes rather short, in which case the proposed algorithm deteriorates in performance. The performance is further exacerbated if new patients join the health program and existing enrolled patients leave the program each day. Some other previous works are able to handle both these issues but their algorithms remain computationally very expensive.

To mitigate these issues, my research proposes the *Streaming Bandits* framework [Mate *et al.*, 2021a] that can accommodate finite horizon settings. It identifies *index decay*, as the key phenomenon behind the performance degradation and proposes an interpolation algorithm to get around the same. This combines the best of both the worlds, achieving a performance quality similar to without being expensive and with a run time similar to [Mate *et al.*, 2020].

4 Conclusion and Future Work

We present a framework based on the Restless Multi-Armed Bandit model, to help address challenges pertaining to planning of limited intervention resources in public health settings. While we address a few challenges, there are still many avenues in this space that need attention. One key challenge is learning the transition dynamics of the patients that our model needs as input. Traditional bandit approaches consider a regret-based framework that learns to balance exploration-exploitation on the fly. However, these approaches may be inadequate as human health workers seem to infer patient characteristics much quicker through a few interactions. Integrating the valuable human input and unlocking the health worker's field expertise through human-in-the-loop approaches is another promising direction of interest.

References

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

[Mate *et al.*, 2020] Aditya Mate, Jackson A Killian, Haifeng Xu, Andrew Perrault, and Milind Tambe. Collapsing bandits and their application to public health interventions. In *Advances in Neural and Information Processing Systems (NeurIPS) 2020*, 2020.

[Mate *et al.*, 2021a] Aditya Mate, Arpita Biswas, Christoph Siebenbrunner, and Milind Tambe. Efficient algorithms for finite horizon and streaming restless multi-armed bandit problems. *arXiv preprint arXiv:2103.04730*, 2021.

[Mate *et al.*, 2021b] Aditya Mate, Andrew Perrault, and Milind Tambe. Risk-aware interventions in public health: Planning with restless multi-armed bandits. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2021.

[Sombabu *et al.*, 2020] Bejjipuram Sombabu, Aditya Mate, D. Manjunath, and Sharayu. Moharir. Whittle index for AoI-aware scheduling. In *COMSNETS*. IEEE, 2020.

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