Restless Bandits visiting Villages: A Preliminary Study on distributing Public Health Services

Biswarup Bhattacharya

University of Southern California

bbhattac@usc.edu

June 20, 2018

ACM COMPASS 2018

Overview

Introduction

Distribution of public health services is a major challenge.

- Health workers $\rightarrow$ spreading awareness of health issues.

Introduction

Distribution of public health services is a major challenge.

- Health workers $\rightarrow$ spreading awareness of health issues.
- A problem of limited resources: There is only an average of 1 health worker per 500 individuals in India [Fan and Anand, 2016].

Introduction

Distribution of public health services is a major challenge.

- Health workers $\rightarrow$ spreading awareness of health issues.
- A problem of limited resources: There is only an average of 1 health worker per 500 individuals in India [Fan and Anand, 2016].
- In India: High infant mortality rate, only 15% of mothers receive antenatal care, mental health issues are prevalent [Venkatashiva Reddy et al., 2013].
- Example: Lack of awareness regarding nutrition, immunization, diseases prevention and treatment, maternity care, and family planning in villages of Bihar (a state in India) [Friedman and Somani, 2002].

Abstract Problem Statement

Given a set of n villages and limited resources, determine a visitation policy so as to reach the most number of affected people, i.e. utilize the limited number of health workers in the most effective fashion.

Contributions

Our main contributions include:

1. Macro-level planning (region-level) $\rightarrow$ p-functional regions problem (PFRP).
2. Micro-level planning (village-level) $\rightarrow$ restless multi-armed bandit (RMAB) model w/ Whittle Index Policy + POMDPs.
3. Addressing the heterogeneity of health problems across villages.

Macro-level Planning

Figure: Administrative boundaries (credits: quickgs.com)

Macro-level Planning

Figure: New "public-health district" boundaries (not an actual simulation)

What is the p-functional Regions Problem (PFRP)?

Introduced by [Duque et al., 2012].

- Aggregate n areal units into p contiguous groups.

- Predefined objective function with a given set of criteria or constraints.
- The objective function can be formulated to minimize the dissimilarity of areal units.

What is the p-functional Regions Problem (PFRP)?

Introduced by [Duque et al., 2012].

- Aggregate n areal units into p contiguous groups.
  - $p =$ number of public health districts desired, and $n =$ total number of villages.
- Predefined objective function with a given set of criteria or constraints.
- The objective function can be formulated to minimize the dissimilarity of areal units.

What is the p-functional Regions Problem (PFRP)?

Introduced by [Duque et al., 2012].

- Aggregate n areal units into p contiguous groups.
  - p = number of public health districts desired, and n = total number of villages.
- Predefined objective function with a given set of criteria or constraints.
- The objective function can be formulated to minimize the dissimilarity of areal units.
  - Maximize the similarity in public health issues facing each areal unit.
  - Smallest areal or spatial unit is a village (or a town).

Optimization Objective

The objective function in our case is as follows:

$$\text{Maximize} \sum_k \sum_i c_{ik} x_{ik} \tag{1}$$

- Index k ($k = 1, \ldots, n$) denotes a village selected as a functional center of the region. Index i ($i = 1, \ldots, n$) denotes a village.
- c_{ik}: Denotes the amount of health-related movement between a village i and a functional center k. We model it as a function of population given the available data.
- x_{ik}: Decision variable indicating whether village i is included within region k (1) or not (0).
- Maximizes the interactions between the villages and the functional center within a region.

Optimization Problem Constraints

- The optimization problem simultaneously performs allocation, identifies functional centers and maintains contiguity of regions.
- Two more decision variables exist in the problem: s_{ik} (sinks, to identify functional centers) and f_{ijk} (flows, to maintain contiguity).
- s_{ik}: Decision variable denoting whether village i is chosen to be a functional centre (sink) (1) or not (0).
- f_{ijk}: Decision variable denoting the amount of conceptual flow between villages i and j in region k.
- Constraints are adapted from the well-known PFRP problem and are elaborated in our paper.

Example

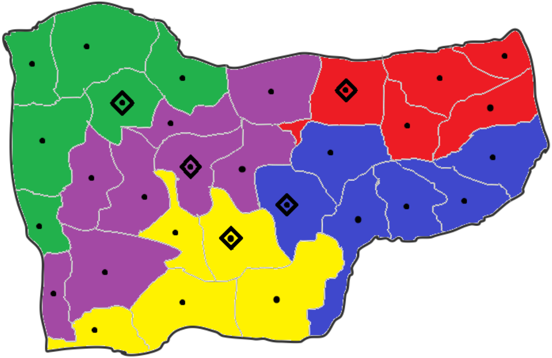

Figure: A solution for $p = 5$

Micro-level Planning Problem Statement

Micro-level planning after macro-level planning for distributing targeted health services.

Given a set of n villages and k ($k << n$) health workers with varied expertise, allocate each health worker to at most one village so as to reach the most number of affected people.

Setup - S, A, O

- Universe of health problems $H = \{H_1, \ldots, H_h\}$.
- For every health problem H_i, the problem is to select k out of n villages to visit.
- Each village has a *hidden health problem intensity* $\mathbf{S}_{H_i} \in \{0, 1, \ldots, n_s - 1\}$. Higher $\mathbf{S}_{H_i}$ implies higher prevalence of that health problem in the village.
- $\mathbf{O}_{H_i} \in \{0, 1, \ldots, n_o - 1\}$.
- Visit = observing *all* h health problems.

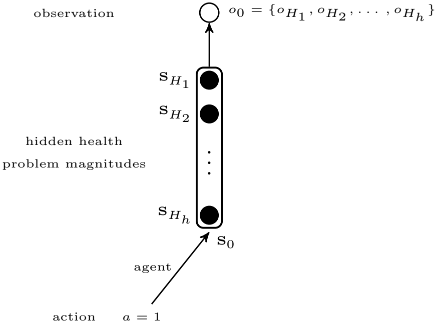

Figure: Hidden health problem magnitudes & observations

Setup - Transitions

Three types of transitions:

- When a health agent of "type" (skill) H_d visits a village:
  1. $\mathbf{S}_{H_d}$ transitions according to a $n_s \times n_s$ transition matrix $\boxed{T^{1,H_d}}$.
  2. All other $\mathbf{S}_{H_i} (i \neq d)$ transition according to $\boxed{T^{1,H_i,gen}}$.
- When a village is not visited:
  3. All $\mathbf{S}$ transition according to $\boxed{T^{0,H_i}}$ (natural rate of deterioration).

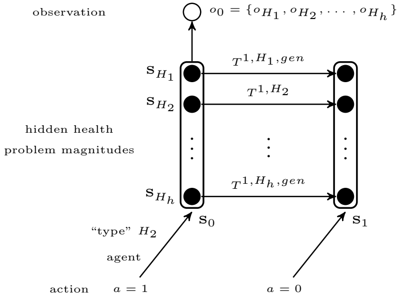

Figure: A two round representation of the Micro-level Planning model thus far

$T^1 \rightarrow$ reduces magnitudes of health problems, $T^0 \rightarrow$ increases the magnitudes.

Micro-level Planning Setup Summary

- When a health agent of “type” H_d visit a village,
  1. She makes an observation regarding every health problem h depending on the current hidden health problem magnitude in that village.
  2. She gets the reward associated with the observation.
  3. Then the hidden health problem magnitude of H_d transitions according to T^{1,H_d}.
  4. Rest of the hidden problem magnitudes H_i, $\forall i \in \{1, \ldots, h\} \setminus \{d\}$ transition according to $T^{1,H_i,gen}$.
- For the villages the health agents do not visit,
  1. They do not have any observation.
  2. They get reward 0.
  3. The hidden health problem magnitude transitions according to T^{0,H_i}, $\forall i$.

Complete Model

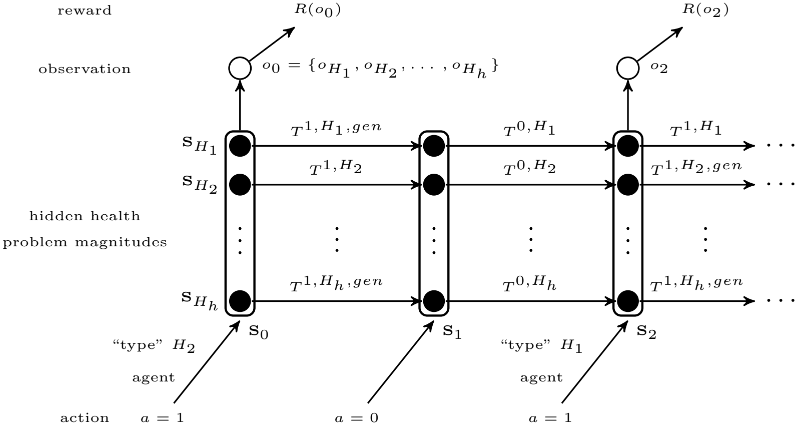

Figure: Complete Model

Restless Bandits

What are restless bandits?

- Multi-armed bandit problem: k out of n arms need to be activated at every round.
- The states of the active arms transition while the states of the passive arms do not change.
- In restless multi-armed bandits (RMABs), the states of the passive arms also transition.

Restless Bandits

What are restless bandits?

- Multi-armed bandit problem: k out of n arms need to be activated at every round.
- The states of the active arms transition while the states of the passive arms do not change.
- In restless multi-armed bandits (RMABs), the states of the passive arms also transition.
- This is close to our setting where the hidden health problem magnitudes (states) can transition even when a village (arm) is not visited (passive arm).
- PSPACE-hard to find the optimal strategy to general RMABs [Papadimitriou and Tsitsiklis, 1999].

Restless Bandits Beliefs

- Exact health states unknown $\rightarrow$ maintaining belief b_{H_i} of each possible state of every possible health problem (H_i) in every village.
- Belief update ("type" H_d health agent when $a = 1$):

$$b'_{H_i}(s') = \begin{cases} \eta_1 \sum_{s \in \mathbf{S}_{H_i}} b_{H_i}(s)\ O^{H_i}_{so}\ T^{1,H_i}_{ss'}, & a = 1, i = d \\ \eta_2 \sum_{s \in \mathbf{S}_{H_i}} b_{H_i}(s)\ O^{H_i}_{so}\ T^{1,H_i,gen}_{ss'}, & a = 1, i \neq d \\ \sum_{s \in \mathbf{S}_{H_i}} b_{H_i}(s)\ T^{0,H_i}_{ss'}, & a = 0. \end{cases} \tag{2}$$

Index Policies

- To solve general RMABs, index policies are used which assigns a value to each arm to measure how rewarding it is to activate an arm at every stage.
- In literature, **Whittle index** is used for RMABs [Whittle, 1988].

Whittle Index Policy

The policy is to activate the arms with the k highest Whittle Indices.

- Concept: The amount of subsidy that needs to be provided to every arm which would make passive action optimal for the current state.
- Larger subsidy $\rightarrow$ larger gap between active action (activate) and passive action $\rightarrow$ more attractive to activate this arm.

Whittle Index Policy

Formally,

$$I(x) \triangleq \inf_m \{m : V_m(x; a = 0) \geq V_m(x; a = 1)\}$$

- Whittle Index of an arm is the smallest m that would make passive action optimal for current state x.
- $V_m(x; a = 0)$ ($V_m(x; a = 1)$) denotes the maximum cumulative reward the player can achieve until the end if he takes passive action (active action) at the first round at the state x with subsidy m.

Whittle Index Policy

Utilizing the belief states:

- When a village is not visited, the immediate reward is the subsidy and there is a β-discounted future reward.

$$V_m(b_{H_i}; a = 0) = m + \beta V_m(b_{H_i, a=0}) \tag{3}$$

Whittle Index Policy

- When a village is visited, there is an expected immediate reward (first term) and there is a β-discounted future reward. $V_m(b^o_{H_i,a=1})$ is the value function at new belief $b^o_{H_i,a=1}$ that is updated from b_{H_i}.

$$V_m(b_{H_i}; a = 1) = \sum_{s \in \mathbf{S}_{H_i}} b_{H_i}(s) \sum_{i=1}^{h} \sum_{o \in \mathbf{O}_{H_i}} O^{H_i}_{so} \, R(o) + \beta \sum_{o \in \mathbf{O}_{H_i}} \sum_{s \in \mathbf{S}_{H_i}} b_{H_i}(s) \, O^{H_i}_{so} \, V_m(b^o_{H_i,a=1}) \tag{4}$$

Whittle Index Policy

- The final value function.

$$V_m(b_{H_i}) = \max\{V_m(b_{H_i}; a = 0), V_m(b_{H_i}; a = 1)\} \qquad (5)$$

- The Whittle Index for belief b_{H_i}.

$$I(b_{H_i}) \triangleq \inf_m\{m : V_m(b_{H_i}; a = 0) \geq V_m(b_{H_i}; a = 1)\} \qquad (6)$$

- The passive action set $\phi_{H_i}(m)$, which is the set of belief states for which passive action ("not visit") is the optimal action given subsidy m.

$$\phi_{H_i}(m) \triangleq \{b_{H_i} : V_m(b_{H_i}; a = 0) \geq V_m(b_{H_i}; a = 1)\} \qquad (7)$$

Numerical Evaluation of Whittle Index

- Indexability requires that for a given state x, its optimal action can never switch from passive action to active action with the increase of m.
- We prove that the RMAB is indexable for $m \subseteq [hR(0) - \beta h \frac{R(n_o-1)-R(0)}{1-\beta}, hR(n_o - 1)]$.
- Given the indexability, the Whittle Index can be found by simply doing a binary search within the range $m \subseteq [hR(0) - \beta h \frac{R(n_o-1)-R(0)}{1-\beta}, hR(n_o - 1)]$.

Planning with POMDPs

- Given the subsidy m, the passive action set $\phi(m)$ can be computed using a POMDP model.
- Every single health problem H_y in every village is modeled as a POMDP – similar to a subset of the RMAB problem setup.
- Given the subsidy m, $\phi(m)$ can be determined by solving a POMDP model which can be set up in a similar way as the problem setup.
- We can combine the POMDP models of every village to form a full POMDP model for a particular health problem.
- Such h complete POMDPs can be created to consider all the health problems.

Experimental Setup

A region of 30 villages (pop. 60,000) in Arwal, a district in Bihar.

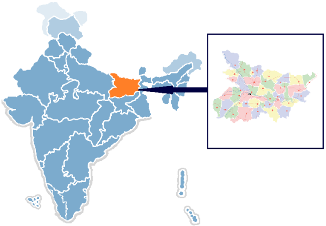

Figure: Bihar (highlighted) in India (image credits: Wikipedia)

- Health agents: ASHAs + ANMs, among others.
- 1 health worker per 1700 individuals.
- 5 Primary Health Centres and 65 Health Subcentres in the whole district.
- Major issues: lack of infant care and ante-natal care, malaria, TB, and leprosy.

Macro-level Planning Data

- Demographic data from the Indian Census 2011 [Government of India, 2011].
- A village = smallest spatial unit.
- Population as the indicator for activity.
- Reasonable proxy for activity (given no other suitable alternatives) as public health is a resource utilized by all people of all ages.

Macro-level Planning Results

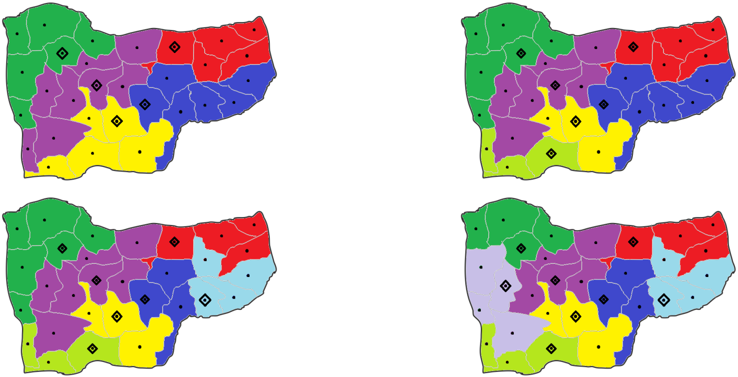

Figure: Illustration of optimal districting of 30 villages for different number of regions ($p = 5, 6, 7, 8$)

Micro-level Planning

- Planning village-by-village visit.
- Now, micro-level planning can be conducted in batches of n/p number of villages (≤ 6 in our case).
- Evaluated in terms of the cumulative reward achieved within the first 30 rounds (discount factor $\beta = 0.9$).

Results

Table: Solution Quality for small-scale problem: $n = 2, k = 1, h = 2, n_s = n_o = 2$

Random	Myopic	POMCP	POMDP	WI
3.280	3.521	3.590	3.717	3.695

- Rewards obtained $\propto$ number of individuals positively affected by the health workers.
- As expected, exact POMDP provides the highest solution quality.
- Whittle Index policy performs reasonably well.

Results

Table: Evaluation for $n = 5, h = 2, n_s = n_o = 2,$ varying k

k	Random	Myopic	POMCP	WI
1	7.441	10.512	11.810	12.124
2	10.431	14.030	15.281	16.425
3	14.129	18.459	18.780	19.294

Table: Evaluation for $n = 5, h = 3, n_s = n_o = 2,$ varying k

k	Random	Myopic	POMCP	WI
1	10.441	16.002	18.111	19.356
2	15.006	20.141	22.164	23.379
3	21.463	23.988	25.665	26.414

Key Takeaways

- ✓ We have presented a **hierarchical model** - PFRP method for macro-planning and a RMAB approach with Whittle Index Policy for micro-planning.
- ✓ A **PFRP** model is used to draw districts based on demographics and health outcome priorities.
- ✓ **RMABs** are effective in modeling the nature of the health problems.
- ✓ **Whittle Index** policies are a suitable alternative to solving exact POMDPs for such kinds of problems.
- ✓ This general structure and setup may also be applicable to various other scenarios of **skill-based service delivery** and **limited resources allocation problems**.

Conclusion

- Similar approaches can be explored by health administrations for planning health policies in the future.
- Important to test and fine-tune this algorithm in a real-world deployment setting.
- Possible directions of future work:
  - Improving the districting model with more “relevant” data.
  - Working on the scalability aspects of the overall model.

Thank You!

Email: bbhattac@usc.edu

Web: https://biswarupb.github.io