Algorithmic Social Intervention

Bryan Wilder
University of Southern California
bwilder@usc.edu

Abstract

Social and behavioral interventions are a critical tool for governments and communities to tackle deep-rooted societal challenges such as homelessness, disease, and poverty. However, real-world interventions are almost always plagued by limited resources and limited data, which creates a computational challenge: how can we use algorithmic techniques to enhance the targeting and delivery of social and behavioral interventions? The goal of my thesis is to provide a unified study of such questions, collectively considered under the name "algorithmic social intervention". This proposal introduces algorithmic social intervention as a distinct area with characteristic technical challenges, presents my published research in the context of these challenges, and outlines open problems for future work. A common technical theme is decision making under uncertainty: how can we find actions which will impact a social system in desirable ways under limitations of knowledge and resources? The primary application area for my work thus far is public health, e.g. HIV or tuberculosis prevention. For instance, I have developed a series of algorithms which optimize social network interventions for HIV prevention. Two of these algorithms have been pilot-tested in collaboration with LA-area service providers for homeless youth, with preliminary results showing substantial improvement over the previously used heuristic. My work also spans other topics in infectious disease prevention and underlying algorithmic questions in robust and risk-aware submodular optimization.

1 Introduction

My research examines how techniques in artificial intelligence (including optimization, machine learning, game theory, and social network analysis) can be used to intervene in social and behavioral systems. Societies around the world face challenges of enormous scale: preventing and treating disease, confronting poverty, and a range of other issues impacting billions of people. In response, governments and communities deploy interventions addressing these problems (e.g., outreach campaigns to enroll patients in treatment or educational programs to raise awareness about preventative strategies). However, these interventions are always subject to limited resources and are deployed under considerable uncertainty; hence, deciding manually on the best way to deploy an intervention is extremely difficult.

Motivated by such challenges, the goal of this thesis is to establish a set of algorithmic techniques which confront underlying challenges in the delivery of social and behavioral interventions (across both public health and other areas) and to field-test these techniques in socially impactful settings. We refer to this domain as *algorithmic social intervention*. Social intervention domains motivate a range of common technical challenges (see Figure 1 and Section 2 for more details). My published work spans all of these areas, though many interesting open problems remain. Specifically, I have studied information gathering [Wilder *et al.*, 2018a; 2018b], optimization under uncertainty [Wilder *et al.*, 2017; Wilder, 2018a; 2018b; Wilder *et al.*, 2018d; 2018c] and adaptive sequential decision making [Wilder *et al.*, 2017; 2018b]. Together with nonprofit agencies, I have empirically evaluated two of the resulting algorithms, with pilot tests showing substantial improvements over status-quo techniques [Yadav *et al.*, 2017; Wilder *et al.*, 2018b].

Specifically, my research thus far has focused on algorithmic approaches to target and enhance interventions in public health settings. One line of work focuses on HIV prevention among homeless youth, where information about HIV is spread through the youths' social network. The challenge is selecting influential peer leaders who will be able to maximize the spread of the resulting diffusion. I have developed a set of algorithms for selecting peer leaders under uncertainty about the structure of the network and how information propagates [Wilder *et al.*, 2017; 2018a; 2018b]. Two of these algorithms have been pilot-tested with LA-area drop in centers serving homeless youth. In these studies, over twice as many youth received HIV prevention messages when the algorithms were used instead of the status-quo heuristic [Wilder *et al.*, 2018b; Yadav *et al.*, 2017]. Another area is infectious disease prevention, where the challenge is to target limited intervention resources (e.g., outreach campaigns to improve treatment uptake) for greatest impact. I developed an algorithm to near-optimally target such interventions, with particular focus on reducing tuberculosis spread in India [Wilder *et*

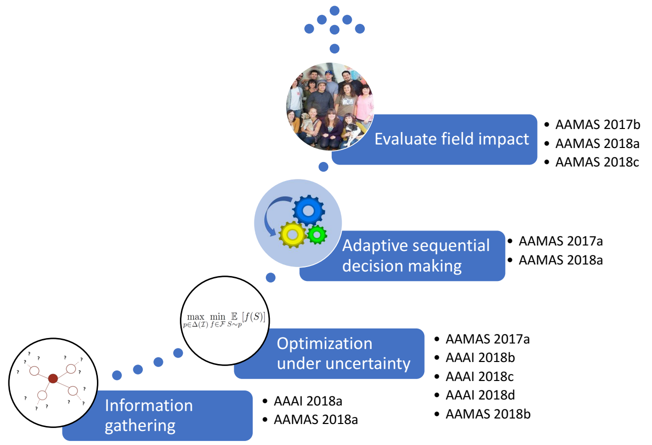

Figure 1: Technical components of algorithmic social intervention and related publications.

al., 2018d]. In simulation, this algorithm averts over 8,000 cases of tuberculosis per year compared to the policy which is currently used.

Underlying these applications are a number of fundamental technical challenges related to decision making under uncertainty. Endemic to public health is a lack of information about the system: where the problems lie, how agents interact, and ultimately what outcome an intervention will have. Similar challenges arise in social and behavioral interventions across numerous contexts. Much of my work formalizes underlying challenges in decision making under uncertainty which are motivated by such applications, develops algorithmic solutions, and proves theoretical guarantees on their performance. The ultimate objective is algorithms that come with both rigorous theoretical analysis and field-tested practical performance. Towards this end, I have developed algorithms for *robust* [Wilder, 2018a] and *risk-averse* [Wilder, 2018b] submodular optimization. Submodularity formalizes a natural diminishing returns property which occurs across many domains (including the HIV and tuberculosis prevention settings), making submodular optimization under uncertainty an important and natural algorithmic challenge.

2 Algorithmic Social Intervention

The goal for this thesis is to establish a unified study of algorithmic social intervention: computational approaches for optimally targeting and enhancing social and behavioral interventions to achieve policy or community-level goals. The aim is to bridge algorithm design, optimization, and machine learning with practice, field deployments, and social impact. Relevant domains are often characterized by common goals and challenges, including:

- Interventions are delivered in a preexisting social context composed of many agents with their own goals and behaviors.
- Agents' behaviors are not totally determined by the intervention: particular incentives, services, or rules may be introduced, but then agents make their own decisions.
- Agents are not perfectly rational, requiring the use of models from the social and behavioral sciences.
- There are many unknowns: agent interactions are complex and are not fully specified by the available data.
- Applications often focus on vulnerable populations.

Figure 1 divides the underlying technical challenges of such domains into several stages. Each stage also lists associated publications. The first stage is *information gathering*. Here, the challenge is to acquire the data needed to optimize the intervention in an efficient manner. For instance, in a social network intervention, it may be necessary to minimize the number of nodes who are surveyed to obtain edges because gathering such data can be extremely time-intensive. The second stage is *optimization under uncertainty*. Since the available data is rarely enough to fully specify the objective function, methods such as robust, stochastic, or risk-aware optimization are necessary. The third stage is *adaptive sequential decision making*. Once an intervention is in progress, the algorithmic system has the opportunity to interact with the world, observe the consequences of its decisions, and adjust accordingly. Lastly, a critical part of algorithmic social intervention is to *evaluate field impact*. While typical means of assessment (theoretical analysis, simulation experiments) are important tools, it is critical to validate the algorithm in a field experiment, ideally in comparison to alternate algorithms.

References

[Wilder *et al.*, 2017] Bryan Wilder, Amulya Yadav, Nicole Immorlica, Eric Rice, and Milind Tambe. Uncharted but not uninfluenced: Influence maximization with an uncertain network. In *AAMAS*, 2017.

[Wilder *et al.*, 2018a] Bryan Wilder, Nicole Immorlica, Eric Rice, and Milind Tambe. Maximizing influence in an unknown social network. In *AAAI*, 2018.

[Wilder *et al.*, 2018b] Bryan Wilder, Laura Onasch-Vera, Juliana Hudson, Jose Luna, Nicole Wilson, Robin Petering, Darlene Woo, Milind Tambe, and Eric Rice. End-to-end influence maximization in the field. In *AAMAS*, 2018.

[Wilder *et al.*, 2018c] Bryan Wilder, Han Ching Ou, Kayla de la Haye, and Milind Tambe. Optimizing network structure for preventative health. In *AAMAS*, 2018.

[Wilder *et al.*, 2018d] Bryan Wilder, Sze-Chuan Suen, and Milind Tambe. Preventing infectious disease in dynamic populations under uncertainty. In *AAAI*, 2018.

[Wilder, 2018a] Bryan Wilder. Equilibrium computation and robust optimization in zero sum games with submodular structure. In *AAAI*, 2018.

[Wilder, 2018b] Bryan Wilder. Risk-sensitive submodular optimization. In *AAAI*, 2018.

[Yadav *et al.*, 2017] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. Influence maximization in the field: The arduous journey from emerging to deployed application. In *AAMAS*, 2017.