

AI at the Margins: Data, Decisions, and Inclusive Social Impact

Bryan Wilder

School of Engineering and Applied Sciences, Harvard University
bwilder@g.harvard.edu

1 Introduction

Artificial intelligence holds tremendous promise to improve human well-being. However, AI techniques are typically developed for the benefit of those with access to technological and financial resources. A critical but understudied question is how AI can benefit marginalized communities who lack such resources. Governments and communities worldwide use a range of interventions to tackle social problems such as homelessness and disease, improving access to opportunity for underserved populations. My research develops machine learning and optimization methods to empower such interventions, which are almost always deployed with limited resources and limited information. Maximizing impact in this context requires algorithmic approaches which span the full pipeline from data to decisions. Computationally, we aim to target scarce resources for greatest effect under substantial uncertainty. The ultimate goal is to include communities worldwide in the benefits of AI progress.

2 Overview of Contributions

My dissertation presents a set of both technical and application-oriented contributions towards this goal. Much of my work has been motivated by a set of application domains, detailed below. However, these domains have inspired a range of general computational challenges at the intersection of optimization and machine learning. This virtuous cycle between field deployment and technical progress lies at the core of my work.

2.1 Motivating Applications and Field Results

One key product of my research is an intervention for HIV prevention among homeless youth which is currently being field tested in collaboration with Los Angeles community organizations. The algorithms that I developed help social workers to identify peer leaders: homeless youth who are well-positioned to spread information about HIV awareness through the youth's social network. The goal is to recruit peer leaders who will influence as many others as possible. The key challenge is that both the social network and the underlying diffusion process are unknown, while the bulk of existing work on influence maximization assumes perfect information. Preliminary results from real-world field trials show that the new methods produce a substantial improvement in

the efficacy of the intervention compared to status quo techniques [Yadav *et al.*, 2017; Wilder *et al.*, 2018b].

In recent work, I have also developed systems which combine machine learning and optimization to improve treatment adherence among tuberculosis patients in India [Wilder *et al.*, 2019; Killian *et al.*, 2019]. Here, the goal is to use a patient's recent adherence pattern combined with demographic data in order to predict which patients are likely to miss doses and optimally target interventions by local health workers. The predictive task is quite difficult due to noisy data and inherent stochasticity: no method will perfectly forecast patients' behavior in advance. Accordingly, machine learning models must be trained specifically to focus on what is important for targeting interventions, extracting as much signal as possible for the optimization problem at hand. I am working in collaboration with NGO and governmental partners to translate these computational methods into a deployable system.

2.2 Technical Contributions

Both applications leverage a series of general technical contributions at the intersection of optimization and machine learning, which enable improved decision making under uncertainty in noisy and data-poor environments. These technical contributions can be roughly arranged along an axis describing how much data is available in any given domain, since this often determines the appropriate technical approach.

When very little data is available, no method will be able to accurately resolve a ground truth model for the impact of potential interventions. Accordingly, decisions must account for uncertainty; formally, we need to solve an optimization problem which depends on unknown parameters. To this end, I have developed methods for *robust*, *risk-averse*, and *stochastic optimization*. The goal is to find decisions (e.g., which youth to select as peer leaders in a social network) that will perform well regardless of the unknown parameters. My primary contribution here is a set of algorithms for submodular optimization under uncertainty (which includes the HIV prevention domain, along with many others). These algorithms efficiently handle problems for which prior techniques would require exponential time ([Wilder, 2018a]), or else do not apply at all ([Wilder, 2018b; Wilder *et al.*, 2018c; Staib *et al.*, 2019]). Beyond HIV prevention, I have used these underlying ideas to tackle additional domains including

obesity prevention [Wilder *et al.*, 2018c] and population-level disease control [Wilder *et al.*, 2018d].

The next case is when little data is initially available, but additional information may be gathered to inform decision making. Gathering data is typically a costly and time-consuming process; it is important to strategically acquire the information which is most useful for informing optimal interventions. I have developed methods for this problem in the context of social network interventions, where the challenge is to survey a small number of population members in order to reconstruct their underlying social network and target influential nodes with a message [Wilder *et al.*, 2018a; Wilder *et al.*, 2018b]. In the HIV prevention domain, these methods substantially reduce the cost to deploy an intervention by minimizing data collection requirements with little loss in solution quality.

Finally, in some cases sufficient data is available for machine learning techniques to recover estimates of the unknowns. However, particularly in the complex and noisy domains which are common to social applications, no model can recover the ground truth with high fidelity. All models make mistakes, and the key is to find a model which nevertheless induces the correct decisions. I have developed methods to integrate combinatorial optimization problems into the training of machine learning models, allowing the training process to automatically focus on the quantities which are most relevant for decision making [Wilder *et al.*, 2019]. The key technical contribution is *differentiable* solvers for common classes of combinatorial problems, allowing models such as neural networks to be trained end-to-end with the optimization problem via gradient descent. I have leveraged these techniques to develop better targeting for interventions in the tuberculosis domain [Killian *et al.*, 2019].

My research also contains another category of technical contributions: field deployment. Creating practical impact is not just a matter of rolling out a general-purpose algorithm. Each domain inevitably comes with new technical challenges which much be addressed. Pilot tests and engagement with domain experts reveal such challenges, which oftentimes inspire new research problems and models [Yadav *et al.*, 2017; Wilder *et al.*, 2018b]. Closing this loop between research and deployment, where engagement with the field provides the inspiration for new computational questions, is a core part of my approach to research.

3 Future Work

Developing AI to benefit marginalized communities presents a wide array of research challenges, with many remaining open questions. One promising area is at the interface of optimization and machine learning. For example, how can we integrate more complex decision problems, e.g., those with harder combinatorial constraints, into machine learning training? This will be a key part of leveraging data for maximum impact in domains where both prediction and optimization are individually difficult. Another important set of concerns relates to issues of fairness and ethics. For instance, how can we ensure that the benefits of an intervention are evenly distributed, without unduly favoring subgroups that are already

advantaged? Both areas represent important topics for future research. Ultimately, creating social impact will require a true integration of fundamental computational research with the on-the-ground needs of communities worldwide.

References

- [Killian *et al.*, 2019] Jackson A Killian, Bryan Wilder, Amit Sharma, Vinod Choudhary, Bistra Dilkina, and Milind Tambe. Learning to prescribe interventions for tuberculosis patients using digital adherence data. *arXiv preprint arXiv:1902.01506*, 2019.
- [Staib *et al.*, 2019] Matthew Staib, Bryan Wilder, and Stefanie Jegelka. Distributionally robust submodular maximization. In *AISTATS*, 2019.
- [Wilder *et al.*, 2018a] Bryan Wilder, Nicole Immerlica, 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 *et al.*, 2019] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *AAAI*, 2019.
- [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.