

Artificial Intelligence for Social Impact: Learning and Planning in the Data-to-Deployment Pipeline

Andrew Perrault, Fei Fang, Arunesh Sinha, Milind Tambe

■ *With the maturing of artificial intelligence (AI) and multiagent systems research, we have a tremendous opportunity to direct these advances toward addressing complex societal problems. In pursuit of this goal of AI for social impact, we as AI researchers must go beyond improvements in computational methodology; it is important to step out in the field to demonstrate social impact. To this end, we focus on the problems of public safety and security, wildlife conservation, and public health in low-resource communities, and present research advances in multiagent systems to address one key cross-cutting challenge: how to effectively deploy our limited intervention resources in these problem domains. We present case studies from our deployments around the world as well as lessons learned that we hope are of use to researchers who are interested in AI for social impact. In pushing this research agenda, we believe AI can indeed play an important role in fighting social injustice and improving society.*

The maturing of artificial intelligence (AI) and multiagent systems research has created a tremendous opportunity to direct these advances toward addressing complex societal problems. In pursuit of this goal of AI for social impact, we as AI researchers must go beyond improvements in computational methodology; it is important to step out in the field to demonstrate social impact. To this end, we focus on three application areas: public safety and security, wildlife conservation, and public health in low-resource communities. In addition to case studies from these areas, we distill our experience into lessons learned and we hope that researchers reading this article may find them useful.

Viewing these societal problems through the lens of multiagent systems, we summarize the goals of our research program as optimizing limited intervention resources when

interacting with other agents. The type of multi-agent interaction varies widely: it might be competitive, where agents are actively trying to achieve different and often conflicting goals, or it might be a process of information spread where the agents do not have explicit goals and just passively react to their surroundings. Our overall research goal is to intervene in this multiagent interaction: to help one of the agents to achieve a desirable social objective. Toward this goal, we develop multiagent system models for the problems, such as game-theoretic models, allowing us to reason about how to maximize our limited intervention resources.

To intervene effectively, we need to understand the details of the interaction and the motivations of the different agents. However, not all elements of the interaction are known. Some elements are partially known through an often incomplete or biased dataset of observations and some are entirely unknown, requiring expert input. In the case where information gathering is time-consuming and costly, we often need to exploit available data to better understand the key latent elements and make more-informed decisions.

Addressing these problems thus requires research advances in several subareas connected to multiagent systems reasoning. For example, new machine-learning models are needed to analyze the data and understand the concealed aspects of the problem. Scalable optimization techniques are needed to design interventions for real-world problem instances.

We take a data-to-deployment approach to AI for social impact research. It begins with immersion, where we seek to understand the problem from the perspective of the decision-making agent, and ends with a field-test, where we validate our modeling approach and algorithms. The data-to-deployment approach is critical because it invites us to refine our models and algorithms to enable direct social impact.

This article summarizes 12 years of work in AI for social impact applied to problems of public safety and security, conservation security, and public health. We provide an overview of this research: our overall research goals, the approach we have found to be successful across domains and objectives, and a history of the projects we've undertaken and their impacts.

The remainder of the article is structured as follows. We begin by defining AI for social impact. We then outline our solution approach: the data-to-deployment pipeline. Next, we discuss specific projects in public safety and security, conservation security, and public health, and the impact these projects have had. We conclude with lessons learned, and a summary.

Defining AI for Social Impact

We find it useful to provide a rough definition of AI for social impact as a subdiscipline within AI.

First, measurable societal impact should be a first-class citizen of this area of research. While a great deal of AI work can be socially beneficial, new research often has no social impact until many years later, when it is refined into a widely usable tool. In the development of computational methodologies, it is often unnecessary to think directly about the end-product — expanding our knowledge and capabilities is a sufficient objective, and rightly so. In thinking about AI for social impact, demonstrating social impact is a key objective. Second, the research primarily focuses on vulnerable groups — for example, the disadvantaged or the endangered — who lack resources to commission beneficial AI research. Third, the research area tended to have not greatly benefitted from AI research in the past. Certain problems are of great direct interest, either commercially or to governments, and as such, have been well-funded throughout the history of AI. AI for social impact focuses on research that would not otherwise be performed if it lacked its impact focus.

AI for social impact work delivers value to the AI community as a whole by providing new problem models; by introducing new contexts to evaluate existing algorithms; and by raising complexities that challenge abstractions, which often motivates extensions to existing techniques. Because AI for social impact work requires extra effort, it requires extra considerations when evaluating its contributions. This is reflected in the Association for the Advancement of Artificial Intelligence 2019 conference and its 2020 AI for Social Impact Track Call for Papers,¹ which states three key aspects where AI for social impact requires more effort than AI that focuses purely on algorithmic improvement. First, data collection may be costly and time-consuming; second, problem modeling may require significant collaborations with domain experts; and third, evaluating social impact may require time-consuming and complex field studies. AI for social-impact researchers invest their resources differently to make contributions to problems of great social importance.

Solution Approach: The Data-to-Deployment Pipeline

We characterize our solution approach as the data-to-deployment pipeline, which is depicted in figure 1. Our activities at each stage of the pipeline are described in the following subsections.

Immersion

In the immersion stage, we seek to gather the available data about the problem and immerse ourselves in the domain. We seek to answer the following questions.

First, who are the agents in the interaction? We want to understand who is making the decisions in the problem. There may be many agents, as

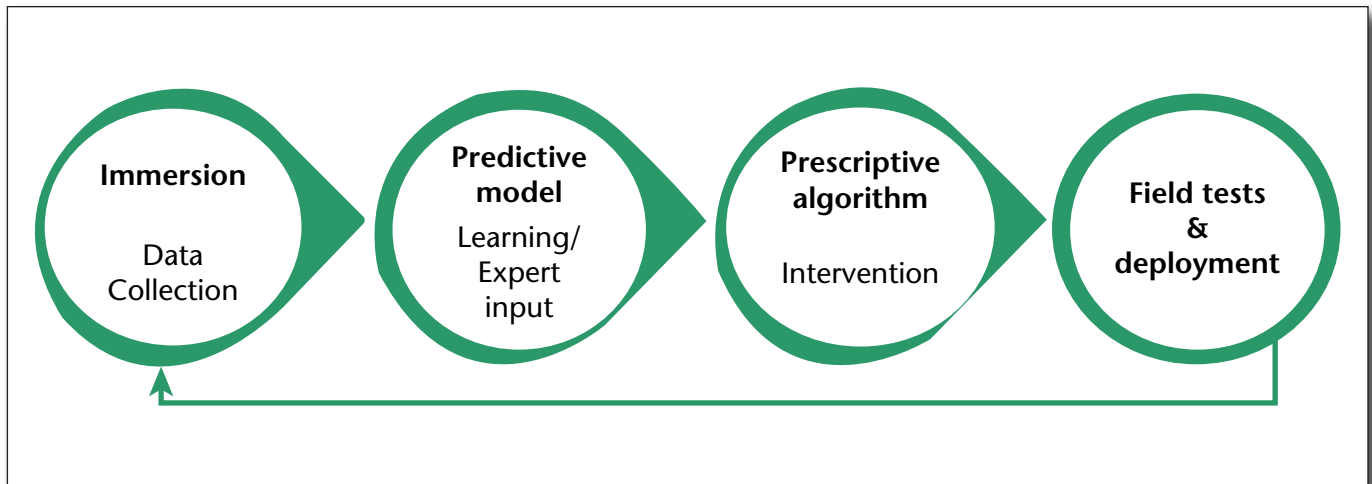


Figure 1. The Data-to-Deployment Pipeline that Describes Our Approach to AI for Social Impact Problems.

in social network interactions, or only two, as in adversarial interactions, such as basic defender-attacker interactions.

Second, what information can agents use to inform their decisions? Addressing this question can be difficult for agents we do not have direct access to. We may make a pessimistic assumption when there is ambiguity: for example, in defender-attacker interactions, we assume that the adversary has access to distributional information about the defender's strategy.

Third, what actions can the agents take, and what impact do they have on the other agents and the environment in which they interact? What is the cost to take each action, and what are the budgets?

These questions may not be answerable directly, but could highlight important latent aspects of the problem that may not be directly observable.

We additionally gather any data that is available from past interactions: the relationships between participants, the effect of actions, the costs or rewards that were accrued, and so forth. During the immersion stage, we often travel to the site of the interaction and talk to the participants directly — this makes it easier to understand the perspective on the ground. We return to the interaction location in the final stage to analyze the impact of the intervention.

Predictive Model

From the immersion stage, we understand the information flow of the interaction and what latent (unobserved) information is critical to defining the interaction. In the predictive modeling stage, we develop a strategy for handling this latent information. A common technique is to build a model that, given the data, makes predictions about high-risk versus low-risk cases, for example,

areas that animal poachers may target, or other classes of relevance.

Prescriptive Algorithm

The output of the predictive model reveals the latent state of the problem that is required to optimize our objective. In this stage, game-theoretic reasoning or multiagent systems reasoning may be used. It is often the case that an optimization problem must be solved, and this may raise computational issues.

Field-Tests and Deployment

Because we take an end-to-end perspective, we must field-test our solutions and compare them to the existing approach. The model we develop is necessarily a simplification of reality, and thus, field-testing is the only way to confirm that we have accomplished our intended goal. This stage relates to the immersion stage, as we return to the field to evaluate our proposed solution and potentially iterate through the design process.

Public Safety and Security

Our research program began in the domain of public safety and security. Motivated by the striking and tragic incidents of terrorist attacks in many parts of the world in the 2000s, we initiated a study of intelligent approaches to thwart attacks on public infrastructure and protect human life. We provide a brief overview of our work in this area. See Sinha et al. (2018) for a comprehensive survey.

Assistant for Randomized Monitoring Over Routes: Security at Los Angeles Airport (2007)

Our work on patrolling the Los Angeles Airport (LAX) was described in Pita et al. (2009). We include it for completeness, as it was the application that inspired this line of research. The terminals of

Stackelberg Security Game Model and Equilibrium

A Stackelberg Security Game is a game played between two players: the defender and the adversary. The defender's task is to protect T targets using $K \ll T$ resources. If the adversary attacks target i , and i is protected by the defender, the defender gets $U_d^c(i)$ and the adversary receives $U_a^c(i)$. Conversely, if not protected, the defender gets $U_d^u(i)$ and the adversary receives $U_a^u(i)$. The defender's randomized allocation (mixed strategy) results in a probability c_i of covering target i , which results in expected utility $c_i * U_d^c(i) + (1 - c_i) * U_d^u(i)$ and $c_i * U_a^c(i) + (1 - c_i) * U_a^u(i)$ for the defender and adversary, respectively. Given a defender mixed strategy, a best-responding adversary chooses a target to attack that maximizes its expected utility. Informally, the Stackelberg equilibrium is the defender mixed strategy that maximizes the defender expected utility against a best responding adversary. For a game-theoretic analysis of general Stackelberg games, see von Stengel and Zamir (2004).

LAX are patrolled by police to ensure the safety of passengers and the protection of infrastructure. As in most security settings, available patrollers cannot monitor every terminal simultaneously. Thus, the patrolling resources must be allocated intelligently, taking into account the differences among the terminals and the adversary's response to information gained by surveilling the patrols.

We model the problem as a Stackelberg security game (SSG) between the defender and an adversary (Pita et al. 2008). The defender's action is a choice from the various combinations of patrol allocations, and the adversary's action is the choice of which terminal to attack. The game's parameters, such as the value gained by the attacker and lost by the defender in the case of a successful attack, were elicited by extensive consultation with airport safety experts — these were ultimately linked to the numbers of lives potentially lost if such an attack were successful, and we were provided extensive data on passengers at different times of day in different parts of the airport. Solving for the game's equilibrium provides the required intelligent randomized strategy. See the inset for a formal description of SSGs.

The deployment of our system for patrol planning at LAX, named the Assistant for Randomized Monitoring Over Routes (ARMOR),² spurred extensive research activity on SSGs. As far as we know, ARMOR was the first deployed application of game theory for operational security recommendations. The successful deployment was enabled by working closely with police officers on the ground and gaining a deep understanding of the problem.

Evaluating ARMOR was especially challenging due to the (fortunate) rarity of security incidents. However, LAX police observed a significant increase in the number of firearm and drug seizures at LAX in the wake of ARMOR's deployment. While internal evaluations led the police to continue using ARMOR for the next 10 years, we provide a more thorough evaluation

of deployed SSG applications through accessible data in Taylor et al. (2017).

Federal Air Marshal Service (2009)

The ARMOR application, which was featured in many news articles and was mentioned in a US Congressional subcommittee hearing, caught the attention of the US Federal Air Marshal Service. The Federal Air Marshal Service aims to deploy armed air marshals on US flights to protect passengers from dangers such as hijacking. As was the case at LAX, there are not enough marshals to cover every flight, making the problem a natural fit for modeling as an SSG. However, the defender's scheduling problem is considerably more challenging because each marshal's patrol must be a cycle. We were, once again, involved in the entire pipeline from immersion to deployment, which yielded the Intelligent Randomization in Scheduling system (Jain et al. 2010a, 2010b). Intelligent Randomization in Scheduling was evaluated independently by the Transport Security Administration and found to be useful, and it is still in deployment today.

Port Resilience Operational/Tactical Enforcement to Combat Terrorism: Port and Ferry Protection Patrols (2013)

A key mission area of the US Coast Guard (USCG) is protecting ports, waterways, and coastal areas. We built the Port Resilience Operational/Tactical Enforcement to Combat Terrorism (PROTECT)³ system to assist the USCG in achieving this mission. One of the innovative aspects of PROTECT is ferry protection. The USCG deploys patrol boats that escort ferries, which presented new technical challenges because the ferries are mobile and the adversary's strategy space is naturally continuous. Our model was deployed to protect ferries in New York, Boston, and Houston (Shieh et al. 2012). USCG




Figure 2. Evaluation of Our Game-Theoretic Approach to Rail-Fare-Evasion Screening.

Our model produces significantly more captures, warnings, and violations than the status quo.

publicly released some of the data from the USCG's evaluation of PROTECT, which demonstrated that PROTECT resulted in less predictable patrolling. Furthermore, USCG reported more illicit activities within the port after PROTECT was deployed, even though no additional resources were deployed.

Rail-fare Evasion in Los Angeles

Our work on screening rail-fare evasion is an important demonstration of how the challenges of real-world deployment can motivate research. While rail-fare evasion has a limited social impact, it provided an ideal testbed for evaluating the SSG approach due to a high volume of incidents and direct access to data. We began by designing a set of prescriptive patrols for transit police, as we had done in previous applications. However, when deployed, we noted that patrollers were unable to execute their assigned schedules because they were constantly being interrupted; for example, by a train running late or the need to handle a medical emergency. The feedback from deployment made us rethink our approach, leading to a sequential, Markov decision process-based patrolling model that accounts for execution uncertainty. The revamped model was tested on the Los Angeles subway system over 21 days in 2013 (Delle Fave et al. 2014) in a randomized test. Figure 2 summarizes the results, which demonstrated

that the game-theoretic approach catches significantly more evaders than the status quo.

Airport Threat Screening

One of the more recent areas of focus in public safety and security are *threat screening games*, which are motivated by the problem of screening airport passengers. An adversary disguises themselves as a passenger and times their arrival to minimize the chance of detection (for example, at a period of high-screening activity and many low-risk passengers). The defender has different types of screening resources, for example, metal detectors and advanced imaging, which screen passengers at different rates. Additionally, the defender has access to data about each passenger's risk category (the US Transportation Security Administration constructs these based on factors such as frequency of travel) and the harm caused if the passenger were to be the adversary. The defender's goal is to balance timely screening with minimizing the chance that an adversary can slip through undetected.

Our initial formulation of threat screening games required that the screenee must be screened in the time window they arrive in (that is, the airport will not accept delays due to screening; Brown et al. 2016). In this formulation, the defender's optimization is



Figure 3. Snares Found in the Srepok Wildlife Sanctuary, Cambodia.

The location of the latest field-test of the Protection Assistant for Wildlife Security.



Figure 4. Rangers on Patrol in the Srepok Wildlife Sanctuary, Cambodia.

how to allocate screening resources to each category of screenees while satisfying the timing requirement. Later variations proposed more complex models:

handling uncertainty in passenger arrivals and different screening rates based on the screenee. These models present the largest and hardest instance of SSGs (Xu 2016). Threat screening games have been tested with real-world airport data. They have also been proposed for problems outside of airport screening such as cybersecurity (Schlenker et al. 2017).

Public safety and security continue to present novel challenges as adversaries innovate. Defenders need to be agile, making use of AI tools to reflect the realities of a changing threat environment.

Conservation Security

The successes in public safety and infrastructure security inspired us to consider what we call conservation security domains that also feature limited law enforcement resources. Illegal activities such as poaching, illegal logging, and illegal, unreported, and unregulated fishing can lead to the destruction of ecosystems. For example, the African elephant population declined by thirty percent between 2007 and 2014, primarily due to illegal poaching. To combat such activities, law enforcement sends patrollers as well as more advanced tools, such as aircraft and drones, to areas of interest to detect and deter illegal activities. However, the patrolling resources are even sparser than those in the public safety and security domain. For example, at one point, only 60 rangers were patrolling Murchison Falls National Park in Uganda, which is almost 4,000 square kilometers.

The role of data is dramatically different in conservation security than in the counter-terrorism tasks mentioned earlier. First, there is much more data available. For example, rangers at the Murchison Falls National Park remove more than a thousand snares per year (figures 3 and 4). They record their patrol routes and the locations of snares using the Spatial Monitoring And Reporting Tool,⁴ creating data that can be analyzed. Second, the data are uncertain in multiple ways — for example, rangers may fail to find a snare even if one is present. The central role of data makes the interaction between game theory and machine learning a key aspect of conservation security research. In this section, we describe two conservation security projects that have traversed the data-to-deployment pipeline.

The Protection Assistant for Wildlife Security

The Protection Assistant for Wildlife Security⁵ is our system for predicting poaching threats and planning ranger patrols to combat poaching. The system consists of three modules: a model to predict poaching behavior; a game-theoretic model for coarse-grained patrol optimization; and a fine-grained patrol planner that takes into account detailed terrain information. Each module has gone through several iterations, and we elaborate on the key developments. The Protection Assistant for Wildlife Security is now being integrated into the Spatial Monitoring And Reporting

Tool, which has been adopted by more than 800 protected areas worldwide, including Srepok Wildlife Sanctuary (Figure 4).

In module 1, we aim to leverage the available data to predict the intensities of poaching activities. Initial versions of this model extended the behavioral game-theoretic approach developed in the public-safety setting (Fang et al. 2016), calculating the subjective utility of poachers as a linear combination of feature values of each target. A target is a cell in a 1-km by 1-km grid representing the protected area. The features of a target may include historical and current patrol effort as well as geospatial features such as animal density, land cover, and slope. A label indicates whether poaching activity was found in the corresponding cell at a particular time.

This approach was only partially successful when applied to real-world data in Queen Elizabeth National Park in Uganda. First, there were very few positive examples relative to the size of the park. Second, we did not handle uncertainty in the data arising from a ranger failing to find a snare even if one is present. More recent work uses more sophisticated machine-learning techniques to address these challenges. For example, Gholami et al. (2018) trains a different classifier for each level of patrol effort and combines them in an ensemble, achieving better predictive accuracy as a result.

We performed extensive validation of the learned models. Our first test sent rangers to two areas in Queen Elizabeth National Park predicted to be poaching hotspots that were not frequently patrolled (Kar et al. 2017). The rangers found three sets of snares in a month, outperforming ninety-one percent of historical months. Following that success, we conducted an 8-month field-test where rangers were sent to 27 areas predicted to be either high or low threat by our model. We found that the catch-per-unit effort, that is, the number of snares found per kilometer of walking, was 10 times higher in the regions that were predicted to be high-threat than those predicted to be low-threat. Later experiments in different protected areas confirmed that our model is effective at identifying and predicting poaching hotspots.

In module 2, we build a game-theoretic model of the interaction between the rangers and the poachers and use it to design patrol strategies that maximize the defender's utility (Xu et al. 2017). We treat the learned model from module 1 as a black box that describes the adversary's behavior, taking the proposed patrol effort and target features as inputs and yielding the probability that a snare will be discovered. The resulting optimization problem is to maximize the expected number of snares discovered by the defender subject to the defender's scheduling constraints, namely that the patroller always starts from the patrol post and must return to it at the end of the patrol, and that patrols have limited distance. We solve this model using mixed-integer linear programming.

While module 2 considers coarse-scheduling constraints, the actual patrols often need to satisfy more fine-grained constraints — complex terrain may

make it impossible for rangers to move from one grid cell to another. In module 3, we incorporate terrain information by building a virtual street map of the area and constructing the patrol strategy on this map (Fang et al. 2016). This module was key to the success in a field-test in Malaysia, where multiple signs of human and animal activity were found.

An avenue for future improvement of the Protection Assistant for Wildlife Security is to consider the interaction between the prediction and game-theoretic models. Our recent work in game-focused learning (Perrault et al. 2020) has shown that including a game model in the machine-learning pipeline improves the defender's utility.

Systematic Poacher Detector for Conservation Drones

Drones can be a valuable patrolling tool. They can be equipped with long-wave thermal infrared cameras, allowing them to effectively detect poachers at night when many poachers are active. The video is then transmitted in real time to ranger stations. Drones present three main technical challenges. First, monitoring drone-captured video is tedious. Second, drones cannot directly interdict the poachers and force them to leave the area, therefore, the drones and rangers must be coordinated. Third, drones can display a flashing light, alerting poachers that they are being observed (this signaling capability, if used carefully, can dissuade poaching activity through the threat that a ranger will be dispatched; however, if overused, signals lose credibility and poachers ignore them).

The Systematic Poacher Detector⁶ is designed to tackle the first challenge. It augments conservation drones with the ability to automatically detect humans and animals in near-real time (Bondi et al. 2018). Given historical videos taken by unmanned aerial vehicle systems, we treat each video frame as an image and collect labels (bounding boxes) for any humans or animals. Our deep-learning-based model leverages available computing resources (for example, graphics processing unit laptops, cloud computing) to improve the detection speed of Systematic Poacher Detector in the field. Air Shepherd,⁷ a drone-based conservation program, conducted a real-world test, with promising results (see Figure 5).

To plan the coordination of drones and human patrollers as well as the signaling scheme, we built a Sensor-Empowered Security Game model based on SSGs (Xu et al. 2017). We show that, in the optimal signaling scheme, the drones always send a warning signal when there is a nearby ranger and send a deceptive warning signal with a carefully designed probability when there is no nearby patroller. Simulation results show that well-coordinated deployment and signaling significantly benefits the rangers. This model assumes that drones always detect a poacher when one is present, and we are currently working to extend the model to account for detection uncertainty.




Figure 5. Systematic Poacher Detector Was Able to Detect Humans in a Test Run by Non-Governmental Organization Partner, Air Shepherd.

Public Health

In this section, we describe two major public health projects we have undertaken. The first focuses on spreading information to prevent human immunodeficiency virus (HIV) among homeless youth in Los Angeles. The second aims to improve tuberculosis medication adherence in India.

Preventing the Spread of HIV Among Homeless Youth

Homelessness affects around 2 million youths in the United States annually, eleven percent of whom are infected with HIV, which is 10 times the rate of infection in the general population (Aidala and Sumartojo 2007). Peer-led HIV prevention programs such as Popular Opinion Leader (Kelly et al. 1997) try to spread information about HIV prevention through a social network of homeless youth by identifying peer leaders within the network to champion the message. The traditional strategy for selecting peer leaders is via *degree centrality* — that is, nodes with the highest number of friendships are picked first. Such peer-led programs are highly desirable to agencies working with homeless youth as these youth are often disengaged from traditional health-care settings and are distrustful of adults. Strategically choosing intervention participants is important so that information percolates through their social network in the most efficient way.

We formulate the problem of selecting peer leaders to spread HIV prevention information as influence maximization with uncertain parameters over an uncertain network (see Figure 6). We assume that the underlying process that is spreading information is an independent cascade model (Kimura and Saito 2006) on a graph $G=(V,E)$ and an associated

function $f(v)$, which represents the probability that influence spreads across edge v . We are uncertain about $f(v)$ and want to maximize the number of influenced nodes in a robust way. We show that we can achieve this objective by formulating the problem as a game against nature, where nature chooses f in response to our choice of seeds, then solving it via double oracle (Wilder et al. 2017). This approach yields an equilibrium strategy despite the exponential search space for the players and converges with approximation guarantees.

A further complication that arises in practice is the unavailability of peer leaders that we selected. For instance, a youth may have gotten arrested or gone to stay with relatives. Thus, we instead think about the problem as choosing a set of peer leaders each week for many weeks according to a training budget. In each successive week, we discover which youth were able to participate last time, informing which new youths to invite this week to continue to maximize information spread. The resulting problem can be formulated as a partially observable Markov decision process and solved via partially observable Markov decision-process decomposition, yielding the HEALER algorithm (Yadav et al. 2015).

We performed a pilot field-test of HEALER, comparing it to the most popular baseline of degree centrality. We selected communities of 60 youths at different centers for homeless youth and our collaborators in social work trained 12 of those youths to be peer leaders (Rice et al. 2018). HEALER is significantly more effective at spreading information in these tests — it reaches around seventy-five percent of non-peer leaders, compared with only twenty-five percent for degree centrality (see figure 7). As a result, HEALER is more effective at causing youth to start testing for HIV: around thirty to forty percent of the community began testing, compared with zero percent for degree centrality.

However, despite its greater effectiveness, HEALER incurs higher costs than degree centrality because it requires that the entire social network be surveyed via on-the-ground work by social workers over many weeks. To overcome this obstacle, we develop a variant of HEALER that only surveys the connections among a small subset of youth as seen in figure 8 (Wilder et al. 2018a). This algorithm, CHANGE, performed as well in field-tests as HEALER (see figure 7), while surveying only eighteen percent of the youth in the network — a major cost reduction.

In other work, we have modeled social influence over a network to optimize public health objectives including preventing childhood obesity in the Antelope Valley in Los Angeles (Wilder et al. 2018b) and preventing suicide among college students (Rahmattalabi et al. 2019b).

Ensuring Tuberculosis Medication Adherence

Tuberculosis (TB) is one of the top 10 causes of death worldwide, and is the deadliest infectious disease;



Figure 6. Social Workers Educate Peer Leaders About HIV Prevention.

This is information that the peer leader is to disseminate in their social network.

last year alone, approximately 10 million people across the globe were infected with TB, leading to 1.8 million deaths. The prevalence of TB is partly attributable to its disproportionate effect on the world's global south where the poor have extremely limited access to healthcare, clean living conditions, and education, which all contribute to the spread of the disease. Further, multi-drug-resistant strains of TB, which are far more expensive and difficult to treat than drug-susceptible TB strains, have taken hold in the world's global south.

The prevalence of TB is caused in part by nonadherence to medication, resulting in a greater risk of death, reinfection, and contraction of drug-resistant TB. To combat nonadherence, the World Health Organization recommends directly observed treatment, in which a health worker confirms that a patient is consuming the required medication daily by observing the patient taking the medication. However, requiring patients to travel to the directly observed treatment facility imposes a financial burden and potential social stigma due to public fear of the disease. Such barriers contribute to patients dropping from treatment, making TB eradication difficult. Thus, digital adherence technologies (DATs), which give patients flexible means to prove adherence, have gained global popularity (Subbaraman et al. 2018).

DATs allow patients to be observed consuming their medication electronically, for example via two-way text messaging, video capture, electronic pillboxes, or toll-free phone calls. Health workers can then view real-time patient adherence on a dashboard such as the one seen in figure 10. In addition to improving patient flexibility and privacy, the dashboard enables health workers to triage patients and focus their limited resources on the highest-risk patients.

Our objective is to use the longitudinal data collected by DATs to help health workers better triage TB patients and deliver interventions to boost the overall adherence of their cohorts (Killian et al. 2019). At first glance, the problem of predicting whom to target for an intervention appears to be a simple supervised machine-learning problem. Given data about a patient's medication adherence, one can train a machine-learning model to predict whether they will miss medication doses in the future. However, such a model ignores the concurrent interventions from health workers as the data were collected, and can lead to incorrect prioritization decisions even when it is highly accurate. For instance, we might observe that missed doses are followed by a period of medication adherence: this does not mean that people with missed doses are more likely to take medication

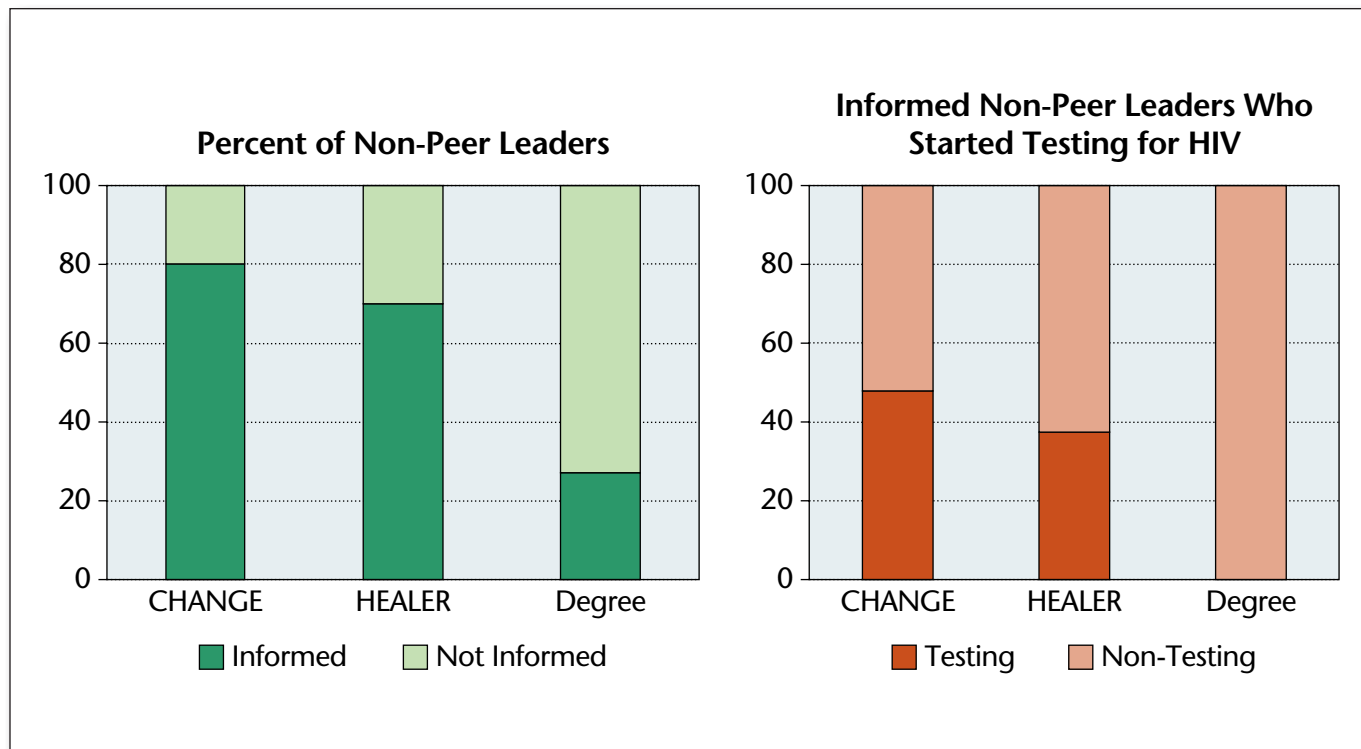


Figure 7. Three Approaches to Selecting Peer Leaders to Spread Information About HIV Prevention.

HEALER, our algorithm that uses network structure to select nodes, outperforms degree centrality (Degree) in both the percent of non-peer leaders reached and the percent of non-peer leaders who began testing for HIV. CHANGE, which uses only partial network information, performs as well as HEALER at a lower surveying cost.

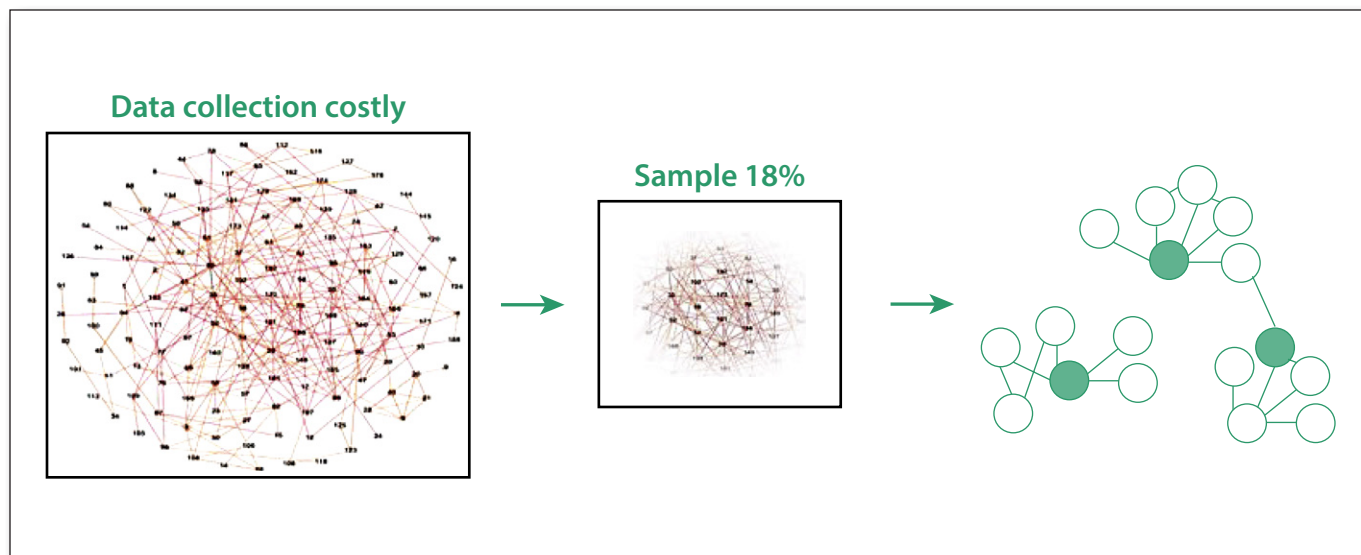


Figure 8. We Decide How to Spread HIV Prevention Information Across a Network by Sampling a Small Number of Edges.

but, most likely, that there was an intervention by a health worker after which the patient restarted their medication.

We introduce a general approach for learning from adherence data with unobserved interventions,

based on domain knowledge of the intervention rules applied by health workers. Using data from the DAT operated by the City TB Office of Mumbai (see figure 9), we show that our approach enables health workers to identify twenty-one percent

more high-risk patients and catch seventy-six percent more missed doses than the currently used heuristics.

We can further improve outcomes by using an end-to-end, decision-focused learning approach (Wilder et al. 2019). Such approaches focus on making predictions that induce good downstream decisions — such as choosing patients for interventions — rather than making perfectly accurate predictions about adherence. In our setup, this approach tunes our system to be more accurate among those patients who could benefit from intervention, rather than being equally accurate across all patients. We find that such a classifier improves the number of successful interventions by approximately fifteen percent compared with a non-decision-focused approach, despite being less accurate about future medication adherence.

Lessons Learned

Based on the experience of the work discussed so far, we state six broad lessons that we have found generally useful. The first two are philosophical (what perspective should we take as AI for social impact researchers), the third is technical, and the remainder relate to the multidisciplinary nature of AI for social impact work.

Take a Data-to-Deployment Perspective

We select projects that can lead directly to real-world deployment in the near future. An academic approach that emphasizes improvements in computational methodology is not necessarily well-suited to achieving this goal — we need to be able to take all the steps from accessing relevant data to deploying prototypes in the field.

Go Out into the Field

Often AI for social impact entails working with vulnerable communities and in remote areas. It is difficult to understand the problems we are trying to solve without consulting the users in the field directly and eliciting crucial details that would not have come to light in the laboratory setting. Additionally, visiting a site allows researchers to understand what technological resources (for example, level of computing power, connectivity) will be available to the intended end-user of the AI solution.

Lack of Data Is the Norm and Needs to Be Addressed in the Project Strategy

It is rarely the case that sufficient data exists in a social impact setting, and developing strategies to address the lack of data is a critical element of our work. For an example project where we apply these strategies, see our project on preventing the spread of HIV among homeless youth.

The first strategy is to make data acquisition part of the deployment plan. If a partner is sufficiently motivated to implement an AI solution, collecting data can energize people working on the ground.




Figure 9. TB Control Center in Mumbai, India.

Collecting data about the existing interaction between agents on the ground is the first step in adapting to an AI approach.

The second strategy is to make data acquisition part of the technical contribution of the project. If data are difficult to acquire, choosing how to collect it can be part of the AI problem (for example, through active learning, preference elicitation, or reinforcement learning). For a solution to be sustainable, the cost of collecting the necessary data must be less than the benefit the solution provides.

The third strategy is to consider sparse data when selecting algorithms. For example, much recent progress in machine learning has focused on cases where there is a large amount of labeled or unlabeled data available. When these conditions are not met, older, statistical approaches may perform better.

The fourth strategy is to consider expert-input or human-subject experiments. In some circumstances, data are so rare, expensive, or sensitive that techniques driven by real-world data are not suitable. This problem arises especially in public-security settings, where attacks can rarely be observed.

AI for Social Impact Work Should be Evaluated Differently Than Other AI Areas

Significant amounts of time and effort must be spent on developing partnerships, modeling, and evaluation to perform research that has a concrete near-term impact. These areas of emphasis require a different approach to evaluation, compared with the one traditionally used at AI conferences.




Figure 10. 99DOTS (Directly Observed Treatment) Electronic Adherence Dashboard Seen by Health Workers for a Given Month.

Missed doses are marked in red, and consumed doses are marked in green.

Build Interdisciplinary Partnerships

AI for social impact work cannot be done without partnerships with researchers in other disciplines who are experts on social impact problems. AI researchers are, by necessity, primarily focused on the problems that arise from the perspective of AI methodology. Thus, if AI is to have a real-world positive impact, it is necessary to leverage expert perspectives on the problems we are trying to address.

Fairness: An Emerging Concern

In research done so far, fairness has been a part of the ethos of partner organizations. As they have been more aware of the challenge of bias in AI systems, questions of fairness have been arising in our research. These issues are quite complicated. While we are currently exploring algorithmic solutions to some of the issue raised (Tsang et al., 2019; Rahmattalabi et al. 2019a), a key question for future investigation is to understand the interaction between domain-specific stakeholder perspectives on fairness and algorithmic approaches.

Summary

Looking to the future, we believe AI is important for improving society and fighting social injustice. To that end, in pushing forward the agenda of AI for social impact, we need to engage in interdisciplinary collaborations and bring the benefits of AI to populations that have not benefited from it. We hope that the case studies we provided and the insights we have gathered are useful.

In many other disciplines, such as human-computer interaction and social work, descriptive work is publishable on its own (for example, Ismail and Kumar, 2019) and may be used as a jumping-off point for intervention design (Fraser and Galinski, 2010). In AI, the descriptive work performed in the immersion stage is a necessary prerequisite for building an AI system, but would not generally be publishable in an AI venue unless paired with the deployment of an intervention.

Notes

1. aaai.org/Conferences/AAAI-20/aaai20specialtrackcall/
2. <https://create.usc.edu/researcher/milind-tambe/projects/armor-assistant-randomized-monitoring-over-routes>
3. <https://www.defensemideanetwork.com/stories/the-port-resiliency-for-operational-tactical-enforcement-to-combat-terrorism-model/2/>
4. smartconservationtools.org
5. <https://sc.cs.cmu.edu/research-detail/102-protection-assistant-for-wildlife-security>
6. <https://dronebelow.com/2018/09/19/drones-help-spot-poachers-in-action/>
7. <https://airshepherd.org/>

References

Aidala, A. A., and Sumartojo, E. 2007. Why Housing? *AIDS and Behavior* 11(2): 1–6. doi.org/10.1007/s10461-007-9302-z.

Bondi, E.; Fang, F.; Hamilton, M.; Kar, D.; Dmello, D.; Choi, J.; Hannaford, R.; Iyer, A.; Joppa, L.; Tambe, M.; and Nevati, R. 2018. SPOT Poachers in Action: Augmenting Conservation Drones with Automatic Detection in Near Real Time. In *Proceedings of the Thirtieth Conference on Innovative Applications*

- of Artificial Intelligence (IAAI 2018). Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Brown, M.; Sinha, A.; Schlenker, A.; and Tambe, M. 2016. One Size Does Not Fit All: A Game-Theoretic Approach for Dynamically and Effectively Screening for Threats. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Delle Fave, F. M.; Jiang, A. X.; Yin, Z.; Zhang, C.; Tambe, M.; Kraus, S.; and Sullivan, J. P. 2014. Game-Theoretic Patrolling With Dynamic Execution Uncertainty and a Case Study on a Real Transit System. *Journal of Artificial Intelligence Research* 50: 321–67. doi.org/10.1613/jair.4317.
- Fang, F.; Nguyen, T. H.; Pickles, R.; Lam, W. Y.; Clements, G. R.; An, B.; Singh, A.; Tambe, M.; and Lemieux, A. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Fraser, M. W., and Galinsky, M. J. 2010. Steps in Intervention Research: Designing and Developing Social Programs. *Research on Social Work Practice* 20(5): 459–66. doi.org/10.1177/1049731509358424.
- Gholami, S.; McCarthy, S.; Dilkina, B.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; Mabonga, J.; and Okello, T. 2018. Adversary Models Account for Imperfect Crime Data: Forecasting and Planning Against Real-World Poachers. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Ismail, A., and Kumar, N. 2019. Empowerment on the Margins: The Online Experiences of Community Health Workers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019*. New York: Association for Computing Machinery. doi/10.1145/3290605.3300329.
- Jain, M.; Kardes, E.; Kiekintveld, C.; Ordóñez, E.; and Tambe, M. Security Games With Arbitrary Schedules: A Branch and Price Approach. 2010a. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, E. 2010b. Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshal Service. *Interfaces* 40(4): 267–90. doi.org/10.1287/inte.1100.0505.
- Kar, D.; Ford, B.; Gholami, S.; Fang, F.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 2017. Cloudy with a Chance of Poaching: Adversary Behavior Modeling and Forecasting with Real-World Poaching Data. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*. Lecture Notes in Computer Science. 1365. Berlin: Springer.
- Kelly, J. A.; Murphy, D. A.; Sikkema, K. J.; McAuliffe, T. L.; Roffman, R. A.; Solomon, L. J.; Winett, R. A.; and Kalichman, S. C. 1997. Randomised, Controlled, Community-Level HIV-Prevention Intervention for Sexual-Risk Behaviour Among Homosexual Men in US Cities. *Lancet* 350(9090): 1500–5. doi.org/10.1016/S0140-6736(97)07439-4.
- Killian, J. A.; Wilder, B.; Sharma, A.; Choudhary, V.; Dilkina, B.; and Tambe, M. 2019. Learning to Prescribe Interventions for Tuberculosis Patients Using Digital Adherence Data. In *Proceedings of the 25th ACM Special Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International Conference on Knowledge Discovery & Data Mining (KDD 2019)*. New York: Association for Computing Machinery. doi.org/10.1145/3292500.3330777.
- Kimura, M., and Saito, K. 2006. Tractable Models for Information Diffusion in Social Networks. In *Knowledge Discovery in Databases: 10th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD 2006)*. Lecture Notes in Computer Science. 1365: pp 259–71. Berlin: Springer.
- Perrault, A.; Wilder, B.; Ewing, E.; Mate, A.; Dilkina, B.; and Tambe, M. 2020. End-to-End Game-Focused Learning of Adversary Behavior in Security Games. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Pita, J.; Jain, M.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2009. Using Game Theory for Los Angeles Airport Security. *AI Magazine* 30(1): 43. doi.org/10.1609/aimag.v30i1.2173.
- Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordóñez, E.; Kraus, S.; and Paruchuri, P. 2008. Deployed ARMOR protection: The Application of a Game-Theoretic Model for Security at the Los Angeles International Airport. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Rahmattalabi, A.; Vayanos, P.; Fulginiti, A.; Rice, E.; Wilder, B.; Yadav, A.; and Tambe, M. Exploring Algorithmic Fairness in Robust Graph Covering Problems. 2019a. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019 (NeurIPS 2019)*, Vancouver, Canada. papers.nips.cc/paper/9707-exploring-algorithmic-fairness-in-robust-graph-covering-problems.pdf.
- Rahmattalabi, A.; Vayanos, P.; Fulginiti, A.; and Tambe, M. 2019b. Robust Peer-Monitoring on Graphs with an Application to Suicide Prevention in Social Networks. In *Proceedings of the Eighteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Rice, E.; Yoshioka-Maxwell, A.; Petering, R.; Onasch-Vera, L.; Craddock, J.; Tambe, M.; Yadav, A.; Wilder, B.; Woo, D.; Winetrobe, H.; and Wilson, N. 2018. Piloting the Use of Artificial Intelligence to Enhance HIV Prevention Interventions for Youth Experiencing Homelessness. *Journal of the Society for Social Work and Research* 9(4): 551–73. doi.org/10.1086/701439.
- Schlenker, A.; Xu, H.; Guirguis, M.; Kiekintveld, C.; Sinha, A.; Tambe, M.; Sonya, S.; Balderas, D.; and Dunstatter, N. 2017. Don't Bury your Head in Warnings: A Game-Theoretic Approach for Intelligent Allocation of Cyber-security Alerts. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*. Freiburg, Germany: IJCAI. doi.org/10.24963/ijcai.2017/54.
- Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: A Deployed Game Theoretic System to Protect the Ports of the United States. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg Security Games: Looking Beyond a

- Decade of Success. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*. Freiburg, Germany: IJCAI. doi.org/10.24963/ijcai.2018/775.
- Subbaraman, R.; de Mondesert, L.; Musiimenta, A.; Pai, M.; Mayer, K. H.; Thomas, B. E.; and Haberer, J. 2018. Digital Adherence Technologies for the Management of Tuberculosis Therapy: Mapping the Landscape and Research Priorities. *BMJ Global Health* 3(5): e001018. doi.org/10.1136/bmjgh-2018-001018.
- Taylor, M.; Kiekintveld, C.; Shieh, E.; Fave, F.; and Tambe, M. 2017. Evaluating Deployed Decision Support Systems for Security: Challenges, Analysis, and Approaches. In *Improving Homeland Security Decisions*. A. Abbas, M. Tambe, and D. Von Winterfeldt, editors. Cambridge, UK: Cambridge University Press. 627–677., doi.org/10.1017/9781316676714.026.
- Tsang, A.; Wilder, B.; Rice, E.; Tambe, M.; and Zick, Y. 2019. Group-Fairness in Influence Maximization. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*. Freiburg, Germany: IJCAI. www.ijcai.org/Proceedings/2019/0831.pdf.
- von Stengel, B., and Zamir, S. 2004. Leadership with Commitment to Mixed Strategies. Computational, Discrete and Applicable Mathematics (CDAM) Research Report, LSE-CDAM-2004-01. London, UK: London School of Economics. www.cdam.lse.ac.uk/Reports/Files/cdam-2004-01.pdf.
- Wilder, B.; Dilkina, B.; and Tambe, M. 2019. Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press. doi.org/10.1609/aaai.v33i01.33011658.
- Wilder, B.; Onasch-Vera, L.; Hudson, J.; Luna, J.; Wilson, N.; Petering, R.; Woo, D.; Tambe, M.; and Rice, E. End-to-End Influence Maximization in the Field. 2018a. In *Proceedings of the Seventeenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Wilder, B.; Ou, H.; de la Haye, K.; and Tambe, M. Optimizing Network Structure for Preventative Health. 2018b. In *Proceedings of the Seventeenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Wilder, B.; Yadav, Y.; Immorlica, N.; Rice, E.; and Tambe, M. 2017. Uncharted but not Uninfluenced: Influence Maximization with an Uncertain Network. In *Proceedings of the Sixteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*. New York: Association for Computing Machinery.
- Xu, H. 2016. The Mysteries of Security Games: Equilibrium Computation Becomes Combinatorial Algorithm Design. In *Proceedings of the Seventeenth ACM Conference on Economics and Computation (EC2016)*. New York: Association for Computing Machinery. doi.org/10.1145/2940716.2940796.
- Xu, H.; Ford, B.; Fang, F.; Dilkina, B.; Plumtpe, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 2017. Optimal Patrol Planning For Green Security Games With Black-Box Attackers. In *Proceedings of the International Conference on Decision and Game Theory for Security (GameSec 2017)*. Lecture Notes in Computer Science. 10575. Berlin, Germany: Springer. doi.org/10.1007/978-3-319-68711-7_24.
- Yadav, A.; Marcolino, L.; Rice, E.; Petering, R.; Winetrobe, H.; Rhoades, H.; Tambe, M.; and Carmichael, H. 2015. Preventing HIV Spread in Homeless Populations Using PSINET. In *Proceedings of the Twenty-Seventh Conference on Innovative Applications of Artificial Intelligence (IAAI-15)*. Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI) Press.
- Andrew Perrault** is a postdoctoral fellow at the Center for Research on Computation and Society at the John A. Paulson School of Engineering and Applied Science at Harvard University. His research is in the field of AI and multiagent systems, focusing on multiagent interactions where information about the agents is not known a priori and must be queried or learned from data.
- Fei Fang** is an assistant professor at the Institute for Software Research in the School of Computer Science at Carnegie Mellon University. Her research lies in the field of AI and multiagent systems, focusing on integrating learning with game theory and contributing to the theme of AI for social good. Her work has won multiple awards, including Distinguished Paper at the International Joint Conferences on Artificial Intelligence-European Conference on Artificial Intelligence 2018, the Innovative Application Award at the Innovative Applications of Artificial Intelligence 2016, and the Outstanding Paper Award in Computational Sustainability Track at the International Joint Conferences on Artificial Intelligence 2015.
- Arunesh Sinha** is an assistant professor at the School of Information Systems at Singapore Management University. His research is at the intersection of security, machine learning, and game theory. His interests lie in the theoretical aspects of multiagent interaction, machine learning, security, and privacy, along with an emphasis on the real-world applicability of theoretical models. His paper was nominated for the Best Application Paper at the International Conference on Autonomous Agents and Multiagent Systems 2016. He was awarded the Bertucci fellowship upon receiving his PhD at Carnegie Mellon University, in appreciation of his novel research.
- Milind Tambe** is the Gordon McKay Professor of Computer Science and director of the Center for Research on Computation and Society at Harvard University; concurrently, he is also director of “AI for Social Good” at Google Research India. Tambe’s research focuses on advancing multiagent systems research for social good. He is the recipient of the International Joint Conferences on Artificial Intelligence John McCarthy Award, the Association for Computing Machinery/Special Interest Group on Artificial Intelligence Autonomous Agents Research Award, and the Association for the Advancement of Artificial Intelligence Robert S. Engelmore Memorial Lecture award. Tambe is a fellow of the Association for Computing Machinery, and of the Association for the Advancement of Artificial Intelligence. The work reported in this article was performed before Tambe’s joining Google.