## Artificial Intelligence and Social Work

This book marries social work and artificial intelligence (AI) to provide an introductory guide for using AI for social good. Following an introductory chapter laying out approaches and ethical principles of using AI for social work interventions, the book describes in detail an intervention to increase the spread of HIV information by using algorithms to determine the key individuals in a social network of homeless youth. Other chapters present interdisciplinary collaborations between AI and social work students, including a chatbot for sexual health information and algorithms to determine who is at higher stress among persons with type 2 diabetes.

For students, academic researchers, industry leaders, and practitioners, these real-life examples from the USC Center for Artificial Intelligence in Society demonstrate how social work and artificial intelligence can be used in tandem for the greater good.

MILIND TAMBE is Helen N. and Emmett H. Jones Professor of Engineering and Founding Co-Director of the Center for Artificial Intelligence in Society at the University of Southern California. He is a fellow of Association for Advancement of Artificial Intelligence (AAAI) and Association for Computing Machinery (ACM), recipient of ACM/SIGAI Autonomous Agents Research Award, Christopher Columbus Fellowship Foundation Homeland security award, INFORMS Wagner prize in Operations Research, IJCAI John McCarthy Award, and others. He has contributed several foundational papers in AI in areas of intelligent agents and computational game theory, which have received best paper awards at major AI conferences including AAMAS, IJCAI, IAAI.

ERIC RICE is Associate Professor and Founding Co-Director of the Center for Artificial Intelligence in Society at the University of Southern California. An expert in community-based research and social network science and theory, he has spent the past several years working to merge social work science and AI to create solutions for major social problems such as homelessness and HIV. Since 2002, he has worked closely with homeless youth providers in Los Angeles and many other communities across the country to develop novel solutions to support young people across the nation who do not have a home, with the goal of ending youth homelessness.

## Artificial Intelligence for Social Good



Artificial intelligence has a diversity of applications, many with social benefits. Books in this series will offer a multidisciplinary perspective on these applications, ranging from social work to health care to environmental sciences. Each book will survey the AI approaches to a particular societal problem and promising research directions, with case study examples.

This is the first book in this series.

# Artificial Intelligence and Social Work

*Edited by*

MILIND TAMBE
*University of Southern California*

ERIC RICE
*University of Southern California*

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

79 Anson Road, #06–04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781108425995
DOI: 10.1017/9781108669016



First published 2018

Printed in the United States of America by Sheridan Books, Inc.

*A catalogue record for this publication is available from the British Library.*

*Library of Congress Cataloging-in-Publication Data*
Names: Tambe, Milind, 1965– editor. | Rice, Eric (Eric R. W.), editor.
Title: Artificial intelligence and social work / edited by Milind Tambe, University of Southern California, Eric Rice, University of Southern California.
Description: Cambridge, United Kingdom ; New York, NY : Cambridge University Press, 2018. | Series: Artificial intelligence for social good |
Includes bibliographical references and index.
Identifiers: LCCN 2018012813 | ISBN 9781108425995 (hardback : alk. paper) | ISBN 9781108444347 (pbk. : alk. paper)
Subjects: LCSH: Social service–Technological innovations | Artificial intelligence. | Social justice.
Classification: LCC HV40 .A79155 2018 | DDC 361.30285/63–dc23
LC record available at https://lccn.loc.gov/2018012813

ISBN 978-1-108-42599-5 Hardback

# Contents

# Contributors

**Nazanin Alipourfard**
University of Southern California

**Elizabeth Bondi**
University of Southern California

**Jacqueline Brixey**
University of Southern California

**Hau Chan**
University of Nebraska–Lincoln

**Jaih Craddock**
University of Southern California

**Rebecca Funke**
University of Southern California

**Shahrzad Gholami**
University of Southern California

**Setareh Nasihati Gilani**
University of Southern California

**David Gray Grant**
Massachusetts Institute of Technology

**Mary Hemler**
University of Southern California

**Rens Hoegen**
University of Southern California

**Nicole Immorlica**
Microsoft Research

**Ali Jalal-Kamali**
University of Southern California

**Albert Jiang**
Trinity University

**Rajni Kumari**
University of Southern California

**Chloe Legendre**
University of Southern California

**Michail Misyrlis**
University of Southern California

**Kien Nguyen**
University of Southern California

**Tanachat Nilanon**
University of Southern California

**Laura Onasch-Vera**
University of Southern California

**Robin Petering**
University of Southern California

**Aida Rahmattalabi**
University of Southern California

**Eric Rice**
University of Southern California

**Orland Roybal**
California State University, Northridge

**Joshua Rusow**
University of Southern California

**Subhasree Sengupta**
University of Southern California

**Emily Sheng**
University of Southern California

**Karan Singla**
University of Southern California

**Ajitesh Srivastava**
University of Southern California

**Milind Tambe**
University of Southern California

**Nazgol Tavabi**
University of Southern California

**Vivek Tiwari**
University of Southern California

**Luan Tran**
University of Southern California

**Mee-Young Um**
University of Southern California

**Lan Wei**
University of Southern California

**Bryan Wilder**
University of Southern California

**Hailey Winetrobe**
University of Southern California

**Darlene Woo**
University of Southern California

**Haifeng Xu**
University of Southern California

**Amulya Yadav**
University of Southern California

**Amanda Yoshioka-Maxwell**
University of Southern California

**Xusen Yin**
University of Southern California

**Kexin Yu**
University of Southern California

**Behnam Zahiri**
University of Southern California

# PART I

# 1

# Merging Social Work Science and Computer Science for Social Good

Eric Rice and Milind Tambe

Artificial Intelligence (AI) is a ubiquitous part of our daily lives. Particularly and inextricably, AI is linked to our online lives. Your every click and post is being mined by increasingly sophisticated algorithms. Whenever you search for something online, AI is there. AI helps us more successfully navigate real-world traffic, not just online traffic. And because it helps us find information quickly or navigate city streets more effectively, it is easy to appreciate how AI can make our lives better sometimes. There are some in the tech world who view AI as a panacea, capable of solving any cognitive task given enough time, computing power, and rich enough data sets. As two authors who live in Los Angeles, where traffic is awful day and night, we certainly will not argue with some of the gains AI can bring to the daily life.

But AI is not always viewed with such a kind eye. Indeed, many in the scientific community as well as the general public are increasingly concerned about the negative impact AI may have on society. Some of these concerns, such as killer robots, may seem like dystopic science fiction fantasies, but weaponized drones guided by AI may not be so far-fetched. There are also genuine concerns over job loss, as some traditional white-collar jobs may be replaced by algorithms that are "happy" to perform dreary organizational tasks for a fraction of the cost of having a human being perform them. Furthermore, many ethicists, policy makers, pundits, and entrepreneurs with high influence in the tech world have begun to worry about the future of society, particularly when AI is applied in arenas where humans have traditionally made decisions, such as medicine and the law. We are also sympathetic to these fears and think that they need to be taken seriously.

This book, however, offers a third perspective. AI can be used to improve society and fight social injustice. From our perspective, the vast majority of the persons who currently benefit from AI are people who live with a certain

amount of privilege. Largely this is due to the fact that much of AI is wrapped up in the high-tech economy. Yet, many of the most impoverished parts of the world still struggle to obtain internet access. The digital footprints of homeless persons, for example, are far smaller than those of us who live in more affluent circumstances.

What we propose in this book is a radical shift in perspective. We dare you to imagine how AI can be used to benefit the disadvantaged. By bringing together social work scientists and computer scientists, we can tackle thorny, seemingly intractable social problems in brand-new ways. Consider this: If problems such as homelessness, HIV, gender discrimination, substance abuse, domestic violence, and suicide were easy to solve, then we would have solved those problems already. Right? But the reality is that many social problems are incredibly complex and require equally complex solutions.

The complexity of the world can rarely be explained by one single discipline. Solving complex problems with one discipline stands even less of a chance. We believe that much of the strength of our work rests in having created an intellectual space in which computer scientists and social work scientists are equal partners who bring unique knowledge and expertise to the table. From AI come new technologies and approaches, and from social work comes a deep understanding of human behavior and how to intervene effectively in society. As such, we see AI for social good as a new transdisciplinary intellectual space.

## Motivations for Our Research

In some respects, this book is an introduction to the work at the University of Southern California (USC) Center for AI in Society (CAIS). This center is one of the first university-based centers focused on AI for Social Good. More importantly it is the first center to deliberately bring together computer science and social work science. CAIS is a joint venture between the Suzanne Dworak-Peck School of Social Work and the Viterbi School of Engineering. Our mission is to conduct research in Artificial Intelligence to help solve the most difficult social problems facing our world.

At CAIS, we seek to bring researchers from around the world to focus on how computer science can be used to solve social problems. This requires an engagement by social work science and computer science. As a field, AI has just begun to consider the role it can have in promoting social good. Similarly, social work has just begun to engage with computer science to enhance the impact of social work science. It is our intention to have the research emerging

from this new intellectual space impact changes in national and international policy, enhance specific social programs, and move toward the creation of a new field of study that merges the technological expertise of engineering with the domain expertise of social work in solving seemingly intractable social problems.

There are many ways one can define social problems, but at CAIS we largely turn to three major sources of inspiration: the Grand Challenges of Social Work, the Grand Challenges of Engineering, and the Sustainable Development Goals of the United Nations. To learn more about these, we encourage you to explore the following web resources:

- Grand Challenges of Social Work (http://aaswsw.org/grand-challenges-initiative/)
- Grand Challenges of Engineering (www.engineeringchallenges.org/)
- United Nations Sustainable Development Goals (www.un.org/sustainabledevelopment/sustainable-development-goals/)
- United Nations Millennium Development Goals (www.un.org/millenniumgoals/environ.shtml)

The Grand Challenges of Social Work is a list of twelve challenges identified by the National Academy of Social Work to inspire social work practitioners and scientists in the twenty-first century. Progress made on any of these challenges will lead to important improvements in human well-being. Some of the challenges that we are particularly motivated by include issues such as ensuring healthy development for all youth, eradicating social isolation, ending homelessness, and harnessing technology for social good. Some of our projects, such as the HIV prevention program for homeless youth that we describe in detail in this book, draw inspiration from multiple challenges. That project in particular is inspired by the desire to harness technology for social good, end homelessness, and empower homeless youth to ensure healthy development for a group of youth who are often marginalized and discarded by mainstream society.

We also draw inspiration from some of the Grand Challenges of Engineering. Like the Grand Challenges of Social Work, the challenges identified by the National Academy of Engineering are for engineers in the field as well as engineering scientists. Two of the challenges we are particularly motivated by are advancing health informatics and securing cyberspace. We have a set of emerging projects that work with administrative data from health providers and social service providers. We are interested in helping build human–machine collaborative decision-making systems that can improve service delivery in both the health and social service fields. Likewise, we have a set of projects

that have focused on public safety and security, both cyber and physical. Some concrete example projects include our collaborations with the US Coast Guard, the Transportation Security Administration, and different police departments. Here we have used AI for enhancing security resource optimization; these algorithms are in daily use for improving randomized patrol routines and allocation at major airports and ports around the United States, including New York, Los Angeles, and Boston. In some cases, such as the US Coast Guard patrols around the Staten Island Ferry in New York, they have radically altered the patrol patterns and resulted in improved effectiveness.

UN Sustainable Development Goals and UN Millennium Development Goals also help anchor the work of the center. We include these goals alongside the Grand Challenges for Social Work because we feel that those Grand Challenges have a domestic flavor, drawing primarily from issues that impact the United States and other developed nations. We doubt this was an intentional oversight, but at CAIS we are interested in issues that may impact the developing world, and emerging economies. We draw obvious inspiration from sustainable development goals such as good health and well-being, equality, no poverty, and decent work and economic growth, as well as millennium development goals such as eradicating extreme poverty, reducing child mortality, promoting gender equality and empowering women, and combating HIV/AIDS, malaria, and other diseases. But we also look to life on land and to ensuring environmental sustainability. We have actively explored the ways of using AI for environmental sustainability and conservation. Concrete ongoing projects include our continued work with nongovernmental organizations (NGOs) for wildlife conservation, including World Wildlife Fund (WWF), Wildlife Conservation Society (WCS), and PANTHERA to implement AI-based techniques for assisting rangers in the field. Here AI algorithms are used for predicting poacher activity, so that patrols can more effectively stop the poachers before they deploy their traps to kill animals. We hope to write a volume dedicated to work on AI for sustainability in the future.

There is a great deal of overlap across the Grand Challenges and Goals from which we draw our inspiration. Both Social Work and the United Nations prioritize fighting poverty, promoting equality, and ensuring healthy development for children. Both Engineering and the United Nations prioritize clean drinking water and improved health. What these challenges and goals share is a focus on improving lives for all people, especially underserved and marginalized populations throughout the United States and internationally. These set of goals provide important new directions for engineers and social workers. It is important to remember that these goals and challenges are not

easy problems to tackle. In fact, a cynic could look at these lists and quickly dismiss most of the topics. If these were easy things to achieve, they would have been achieved already. What we have come to appreciate is that a merging of social work and engineering empowers each to be more effective in tackling these seemingly impossible challenges. We are not offering a panacea, but we are offering a new path. It is perhaps an unlikely path, but we hope that the contributions in this book will do much to demonstrate the potential we have come to see.

## What Is Social Work

Since this book aims to speak to two very different audiences, social work scientists and computer scientists, a brief overview of just what happens in these two disciplines seems warranted. Unfortunately, a few paragraphs on the topics can hardly hope to do each discipline much justice. Moreover, such a short treatment will inevitably reflect our own personal stances within our respective disciplines. That said, we hope these next few paragraphs help you, our readers, to be grounded to some extent in the work of these two fields.

Social work is fundamentally an intervention-driven discipline. The historical roots of social work and its pioneers, particularly Jane Adams, Elizabeth Fry, and Mary Ellen Richmond, were first and foremost in improving the human condition, particularly for the disadvantaged. Over the past 100-plus years, social work has emerged as a practice and a science. Individual social workers work in a variety of settings. The most common are as members of child welfare agencies, community-based or nongovernmental organizations, as community organizers, and in policy-making organizations. What ties this work together is an overarching drive to fight inequality and to help all people live to their full potential. There has been and continues to be a fundamental focus on marginalized and disadvantaged groups within society and on assisting those communities and the individuals who are a part of those groups to thrive.

Social work science focuses on the complex causal forces that impact human life. Central to the discipline is an understanding of human well-being as existing in the nexus of micro-, messo-, and macro-level forces. When social work thinks about micro-level issues, the focus is on the psychology, biology, and even spirituality of individuals. Messo-level action is concerned primarily with the functioning of social networks and local communities. Macro-level issues are conceptualized as larger social forces, such as institutions and social

norms. The science conducted by social work scientists happens at any of these three levels and often cuts across these levels.

Research conducted by social work scientists can take many shapes, but fundamentally there is always a focus on intervention and improving the lives of humans in need. Some of the goals of social work is specifically about the design and testing of new behavioral and/or mental health interventions. Social work scientists tend to want to create change in the world, and many are not happy to let their interventions be sequestered to university-based work. Thus, in recent years, there has begun to be a focus within the social work field on what is called implementation science, that, is the science of bringing evidence-based interventions to scale in communities.

Social work scientists also do a large amount of observational social science research. Much of this work is based on using statistical techniques to analyze survey data collected from a variety of populations, ranging from nationally representative samples to community-based samples of marginalized groups, such as homeless youth. There has been a recent turn toward using administrative data from social welfare and health care service settings as well. In addition, there are also qualitative methods, such as ethnographic work, focus groups, and individual open-ended interviews. But all of this work is in the service of intervention. Almost every article written by a social work scientist includes a discussion of how their finding impact intervention design or social policy change. Social work science is fundamentally a science of action.

## What Is AI

There is no accepted definition of Artificial Intelligence, but generally speaking, we will consider AI to include computer systems (including robots) that perform tasks that require intelligence. From a discipline that the public knew almost nothing about a few decades ago, AI has grown into a discipline that is often in the news (arguably for the wrong reasons), and has led to national research strategies being developed in the United States and the United Kingdom.

Artificial Intelligence began in the late 1950s. Four scientists are generally considered to be the founding fathers of AI: Allen Newell and Herbert Simon of Carnegie Mellon University; John McCarthy of Stanford University, and Marvin Minsky of the Massachusetts Institute of Technology. Since those early years of optimism, AI has gone through its share of ups and downs; it is worthwhile taking a quick tour of these developments and understanding the roots of the discipline. The early years of optimism in AI were based on computer

programs performing tasks that could be considered difficult for humans, such as playing chess, proving theorems, and solving puzzles. Early progress was dominated by use of logics and symbolic systems, heuristic reasoning, and search algorithms. If we are playing chess or proving a theorem, we do not face uncertainty in our reasoning steps: e.g., statements are either true or false; chess moves are deterministic. Accordingly, probabilistic reasoning and uncertainty did not play a major role in the field. Since interacting with the outside world was not a priority in these tasks, a program could reason about a situation in depth and then react.

Early years of optimism also focused on integrated AI systems that could combine sensory perception or computer vision, machine learning to autonomously learn from interactions of the AI system with the world, automated reasoning and planning to act in the world, possibly speech and dialogue in natural language, social interactions with humans and other entities, and so on. Such intelligent agents include robots, but can also be avatars in virtual worlds. This vision of AI involves gathering data from the world and machine learning to build a predictive model of how the world would evolve. However, it does not stop only at making predictions, but also includes planning and reasoning from first principles, i.e., providing prescriptions and taking actions in the world.

This optimistic period had reached its end in the 1990s, when AI research was accused of not living up to its promises, which led to a general cut in funding and interest in AI. Researchers realized that what was easy for humans (e.g., recognizing objects in images, understanding human speech as English words) turned out to be difficult for AI systems. As AI systems started interacting with the real world, there appeared a need for an entirely new set of tools to handle reasoning in the presence of uncertainty and to provide quick reactions to unanticipated events. These difficulties led to new quantitative tools from decision theory, optimization, and game theory being brought into AI. This period also coincided with an emphasis on fast reactive reasoning (rather than purely deep reasoning) in AI, particularly for robots interacting with the world; this was the beginning of the rise of a new AI, with the early symbolic AI approaches being called "Good Old-Fashioned AI" (or GOFAI).

With increasing computing power and the ability to gather large quantities of data, we now have reached yet another new era of AI. Machine learning using deep neural nets has led to significant new successes in tasks such as recognizing objects in images and natural language understanding. While these new approaches to machine learning have not penetrated all of the traditional areas of AI (e.g., planning and reasoning or social reasoning), this new era has brought significant new interest and investment from industry. It has also

led to a mistaken perception among the general public that AI is only about machine learning and prediction, and that unless there are vast quantities of data, AI cannot function. Often not well understood is AI's ability to provide prescriptions, by reasoning through millions of possibilities based on millions of simulations from first principles of how entities may interact in the world. In our work described in Part II of this book, we often focus on this ability to plan interventions.

## The Unlikely Partnership

AI and social work may seem like an unlikely combination or even a clashing set of disciplines. It turns out, however, that there are three major intellectual points of convergence between our individual research agendas, which make this a wonderful, albeit surprising, partnership.The research agendas embrace complexity in social settings. Likewise, both embrace uncertainty. And perhaps most importantly, both value praxis.

Eric Rice's research agenda has been rooted in the idea that the social world is a deeply complex place. The causes of social injustice are complex, and are the results of the intersections of historical oppression, contemporary societal structures, psychological issues, biological issues, and even spiritual issues. Much of his work prior to connecting with that of Milind Tambe focuses on the complexity of social networks and how the web of interconnections that surround people impact human well-being. Likewise, Tambe's research agenda is driven by complex real-world multi-agent system problems, an area of AI focused on social interactions. This area traditionally focuses on employing machine learning, planning, and reasoning in complex social environments involving interactions among humans, as well as between humans and AI (including software or robots).

In their research, both Tambe and Rice must wrestle with uncertainty. From the perspective of computer science modeling, the probabilistic nature of outcomes is par for the course, and techniques such as sequential decision-making and planning under uncertainty that we will discuss later are predicated on not knowing all of what happens in the world. Social work, likewise, has a deep understanding of the uncertainty of the social world. Human beings are not only complicated; they are idiosyncratic. Social work loathes "one size fits all" solutions and embraces understanding of not just what can be known but also what is not known.

But perhaps the greatest point of interconnection between Tambe and Rice's research agenda is a desire to implement solutions. Rice views social work

science as fundamentally a science of intervention. Even when his scholarly articles focus on theory, methodological developments, or observational research, the intent is always to further the well-being of populations in need. Thus, among Rice's published work, there are articles that do not explicitly test new intervention strategies for homeless youth, but nearly all of these articles end with explicit discussions of how that particular piece of research may impact program development or policy change. Likewise, if you look at Tambe's publications, while some are more theoretical than others, focusing on the development of new algorithms, particularly over the past decade, all of this algorithmic work is explicitly focused on trying to assist in the solution of real-world problems.

## New Science for Both Sides

Social work problems and AI may seem far apart, but we argue that, when brought together they engender powerful new science. We argue that AI for social good, or AI and social work, is a new intellectual space, because what has come from our collaborations is unique to both AI and social work. This is not just an issue of applying "out of the box" AI tools (as is the case sometimes when social scientists discover that new analytic techniques are available), nor is it an issue of doing typical social work interventions but with a technology enhancement (as is sometimes the case with mobile health technology projects). What we have created is a new intellectual space that is generative of innovative science from both the computer science and social work science perspective (Figure 1.1).

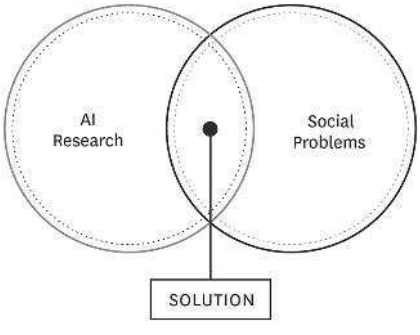


Figure 1.1 The USC Center for Artificial Intelligence in Society's approach to research

We can illustrate this point best in our work on HIV prevention for homeless youth. From the perspective of AI, influence maximization in the context of homeless youth networks introduced several new challenges for which algorithms did not exist. For example, traditional AI influence maximization algorithms, focused on domains such as product advertising, have relied on a "clean" data set of a social network (possibly downloaded from social media) without uncertainty. However, the social network of homeless youth is not downloaded from social media; rather, it is obtained from field observations of social workers, and may be full of uncertainty. Thus, we had to address new research challenges in influence maximization algorithms focused on reasoning with uncertainty in social networks and adapting to availability of new information reducing or adding to such uncertainty. Indeed, this research work has been lauded for its innovation with a number of awards at major AI conferences. Likewise, from the perspective of social work, working on diffusion-based interventions is common, but thinking about diffusion as a formal problem of influence maximization radically shifted the view of networks and actions of actors within those networks. And these innovations were not only significant from a scientific perspective but, as you will see in Chapter 4, made a real impact on the lives of homeless youth.

More generally, there are many new challenges in this interdisciplinary space that do not fit the traditional techniques of either field. As we discussed earlier, the new wave of AI has been driven by large data sets, massive computational power, and hence deep learning; this AI work has often focused on learning predictive models. The type of AI research required for this interdisciplinary space is often devoid of such convenient data sets. Indeed, we are often focused in low-resource communities that have incomplete and uncertain data. Often the massive computational power is in short supply; resources are very limited. Finally, we are often interested in going a step beyond prediction to prescription, to intervention. This style of AI research often calls for partnerships between humans and AI. For example, in our work on HIV prevention for homeless youth, humans are the ones responsible for interacting with homeless youth, whereas the AI programs focus on selecting which youth the social workers should interact with.

Similarly for social work, engaging in formal modeling of social processes in the context of AI allowed us to identify and solve intervention implementation challenges. As we discuss later in the book, peer-led HIV prevention programs have not always worked in all contexts. Part of this failure may have to do with who is picked to be a peer leader. The person who delivers the messages is as important as the messages being delivered. Two approaches to peer leader selection are typically used. Peer leaders are selected either based on attributes that seem to indicate they would possess the capacity to

be a peer leader or by observing a population and selecting the most popular persons. Thinking about peer leader selection within the context of influence maximization, however, shifts the thought process toward the importance of interconnections within the community and the importance of groups of persons working in concert to enact social change. Even static network solutions such as the persons with the most connections to others were abandoned in favor of more dynamic solutions. This approach to intervention development and implementation of social programs is unique to this collaboration between social work and AI.

In addition to the new techniques discussed earlier, this interdisciplinary space also calls for support of new scientific traditions. AI research focuses on improving tools and techniques, and best papers in AI conferences provide technical breakthroughs in models, representations, algorithms, theorems and/or new analysis of algorithms or their theoretical or empirical behaviors. This is fast-paced research, with new breakthroughs coming fast and furious, and often easily evaluated with massive simulations or other evaluations on clean datasets. The focus is "inward," toward the technique. This scientific progress is exciting and obviously important. However, in our interdisciplinary space, it is important to go beyond the technical breakthroughs per se and measure their impact on society. What is required is a careful, patient, and detailed measurement of the pros and cons, benefits and deficits, conducted on the scale of human social interactions, potentially carried out over months. Such detailed understanding of impact of new proposed interventions is absolutely essential, but it is new to traditional AI. The result of such research is not new algorithms or new theorems – the staple of publications in AI – but detailed analysis of impact on society. This is clearly a new kind of science, one that requires nurturing. Similarly for social work, the process of intervention development from acceptability trials, to feasibility trials, to efficacy trials, and finally to effectiveness trials is an incredibly slow-moving process, one that typically takes more than a decade of work with a single new intervention program. We need new, more nimble ways of testing interventions in order to take advantage of the more fast-paced work of computer science. Perhaps something akin to beta-testing is needed for our work to keep pace with the emerging opportunities presented by collaboration with AI.

## The Structure of the Book

This book has two sections: Part I focuses on our ongoing transdisciplinary research specifically around social issues such as homelessness and HIV risk. We focus on how AI can enhance the implementation of HIV prevention

interventions for the homeless youth. Some of this work has been published in a series of papers at various international AI conferences including IAAI 2015, AAMAS 2016, AAMAS 2017 (two separate papers), as well as several workshop papers and *AI Magazine 2016*. We present a version of this work that is more accessible to an interdisciplinary audience rather than a purely AI audience. This work with homeless youth drop-in centers in Los Angeles illustrates that AI algorithms can actually significantly improve spreading of HIV prevention information among homeless youth. The work has received significant number of awards, including AAMAS16 best student paper, AAAI17 best video and best student video award, and a nomination for AAMAS17 best paper. This section also includes an overview of the social context of youth homelessness in the United States, provides an exploration of our first algorithmic solutions to peer-led HIV prevention models, "real-world" field tests of these algorithms with actual homeless youth in Los Angeles, and future thoughts about enhanced algorithmic techniques that rely on less extensive network data collection methods.

Part II of the book features seven aspirational chapters written by graduate students at USC. All of these chapters represent the work of transdisciplinary teams of students that include both social work and computer science graduate students working on a range of difficult social problems using cutting-edge techniques from AI to attempt to create new solutions to wicked social problems. These chapters were the product of a course that we co-taught in the spring of 2017 at USC called "AI for Social Good." The class included students from both social work and computer science, mostly PhD students with a handful of master's students. These chapters showcase many problems that can be tackled by this novel collaboration of AI and social work. The problems covered include interpersonal violence, intimate partner violence, sexual health information dissemination, posttraumatic stress, depression, and homelessness. The AI work includes machine learning techniques such as neural networks, decision trees, and Bayesian networks. There are also chapters that utilize sequential planning under uncertainty, the Uncertain Voter Model, and advances to the k-center problem. There is also an initial foray into natural language processing, with the creation of a Facebook bot to autonomously deliver sexual health information to sexual-minority youth. This section of the book concludes with a chapter on ethics of autonomous agents in the context of social work and public health interventions, written by a CAIS summer fellow who is a doctoral student in philosophy.

We end the book with a glossary of terms. Perhaps the most challenging aspect of AI for social work is that social work science and computer science both have very specific, technical, jargon-filled, academic languages. And

these languages are not at all the same. Over the past few years we have learned that often we use different words to mean the same thing and the same word to mean different things. The word "objective" for an AI researcher is often framed in the context of an optimization problem, and implies a precise mathematical statement, as opposed to the objective that a social scientist may wish to achieve for, say, homeless youth in the real world. The "objective" of a real-world intervention is often something far more nebulous, such as to help prevent the spread of HIV among homeless youth. What a computer scientist calls a feature of a data set a social scientist would call a variable. And perhaps most frustrating, the difference in the word "model" in the context of computer science and social science is so different as to be comical, if it were not so disruptive at times. Because we have experienced the hard growth of trying to become bilingual, we share with you this glossary.