

Introduction: AI for Social Impact

Milind Tambe¹
and Fei Fang² and Bryan Wilder²

¹ Harvard University, USA

² Carnegie Mellon University, USA

`milind_tambe@harvard.edu, {feif,bwilder}@andrew.cmu.edu`

1 What is AI for social impact

We find it useful to begin this chapter by providing an initial definition of *AI for Social Impact*, an emerging sub-discipline of AI. Unlike different areas traditionally defined within AI, e.g., multi-agent systems or reinforcement learning, it is not unified by a specific methodology or a commitment to a particular computational approach.

Instead, this area has key characteristics based on its impact[16]. The first key characteristic of *AI for social impact* is measurable societal impact. Indeed, actual demonstrated social impact is 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; indeed, expanding our knowledge and capabilities is a sufficient objective, and rightly so. In contrast, when thinking about AI for Social Impact, demonstrating social impact is itself a key objective. A second key characteristic of this research area is its primary focus on vulnerable groups, e.g., disadvantaged or endangered, that lack resources to commission beneficial AI research. An additional key characteristic is that it includes research areas or applications that have not greatly benefited 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 provides value to the AI community as a whole by providing new problem models, introducing new contexts to evaluate existing algorithms, and raising complexities that challenge abstractions, which often motivates extensions to existing techniques.

There is also often some difference between more traditional models of methodological AI research and research in AI for social impact. A more traditional project focusing on methodological advances in AI may provide experiments or theorize such advances in the lab, potentially motivated by a variety of challenges – but the key contribution of the work is in its methodological advance. In such a project, we may also provide new research contributions that demonstrate some idea on some simulation or some robots or some artifacts in a lab;

in so doing, we may use real-world data and demonstrate advances in the lab over known benchmarks. This is crucial to AI, as we advance our basic science – models and algorithms. The idea is that such models might eventually influence products and policy.

In contrast, in AI for social impact research, we start with a societal challenge and then attempt to address it by providing the right AI tool. Addressing this societal challenge often must provide some methodological advance, but that may not always be the case; instead, social impact is the key distinguishing feature. Of course, in many cases and perhaps in the most challenging situations of social impact, research is a blend of the two – we have to provide methodological advances in order to show actual impact on society.

Because AI for Social Impact work differs in its emphasis, it requires different considerations when evaluating its contributions. This is reflected in the AAAI AI for Social Impact Track Call for Papers, 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. 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.

In the rest of this chapter and in other sections, we may refer to *AI for social impact* as AI4SI. Note that both *AI for social impact* and *AI for social good* have been used by different groups and authors to get to the concept of a sub-area of AI that is roughly aligned on the characteristics we have outlined above. However, we prefer the term AI for social impact, in part because the word “good” is seen as casting doubt on other research as not being “good”, and because we want to emphasize impact as key characteristic.

2 Why does AI for social impact differ from “traditional AI” – why do we need a different textbook?

These distinctions also manifest in a distinctive set of skills that are required for researchers who aim to develop AI4SI. Standard curricula give students an excellent toolbox for designing and analyzing new algorithms or machine learning models. However, in social impact domains, the questions of what problem to solve and how to evaluate success are often more important (and more difficult) than the computational problem in isolation. Close engagement with domain experts and community members is necessary to shape this process. However, AI researchers need to acquire competencies that will allow them to understand the needs of the communities they aim to work with, translate technical advances into practical interventions, design experiments that measure the impact of such interventions, and rigorously evaluate the results of such experiments. Insights from disciplines such as human-computer interaction (HCI), information and communication technologies in development (ICTD), implementation science in public health and social work, and so on, are therefore essential. Currently, AI

students have little opportunity for systematic instruction in these areas and must learn "on the job". However, we believe that this unique combination of skills and disciplines warrants its own curriculum. Providing the start of such a curriculum is the goal of this book.

The applications covered in this book focus primarily on people and social needs which are *underserved by existing work*. This encompasses work that aims to expand access to opportunity or improve the quality of services offered to marginalized communities, who have been historically neglected by the typically market-driven process of bringing AI from research to application. It also includes efforts to promote public goods, as in environmental conservation or in public health. No individual company may have a sufficient incentive to address such issues because a significant portion of the gain is realized by society as a whole instead of specific individuals (a positive externality). Of course, "public good" domains and underserved groups are closely linked; e.g. the burden of public health or environmental challenges is borne disproportionately by people from marginalized communities. We argue for a greater focus on these domains precisely because existing social or incentive structures do not suffice to meet the needs of those who are impacted. This distinction helps clarify that we are not attempting to strictly delimit what counts as "good" – many businesses contribute to the overall welfare of society, and AI efforts across a wide range of domains not included in this book may have a beneficial impact. However, by definition, there will be no shortage of existing momentum in the field to address such use cases. Rather, we propose that there is an *underexplored* opportunity for AI researchers to have a positive impact through careful engagement with domains like the ones that we cover and that doing so will require a set of skills and values that are not covered in standard AI curricula.

3 History and Trend of Research on AI for Social Impact

There have been scattered research discussions and efforts that focus on using AI to tackle various societal challenges since the field of AI has been established [13, 8]. In the past two decades, there is a significant increase in the interest in AI4SI within the AI research community, and the number of research publications on AI4SI has increased steadily. There are a few catalysts for this increase.

First, we have seen the tremendous success of AI in various tasks. This is due to the development of new AI methods, significantly increased computing capacity, and more publicly available datasets. For example, training carefully designed convolutional neural networks (CNNs) on large datasets such as ImageNet [9] using novel learning algorithms has led to superior performance in computer vision (CV) related tasks such as image classification, object detection, image segmentation, image generation. Advanced pre-trained language models [10] significantly improved the performance in natural language processing (NLP) related tasks such as question answering, translation, and language inference. Search and reinforcement learning (RL)-powered game AIs have led to a superhuman performance in card and board games like poker, Go [7, 21],

and real-time strategy games such as Dota 2 and Starcraft [14, 5, 22]. Many of these advances inspire new ways to tackle various societal challenges we face. For example, advanced CV techniques have been used to identify individual wildlife or marine life from camera traps and photos taken by tourists, which is important to choose appropriate conservation actions [19, 6, 15]. Pre-trained NLP models can be used to identify homeless youth at-risk of substance use by analyzing their answers to surveys and social media posts [11]. RL algorithms are used to improve transportation efficiency by improving vehicle dispatching on ridesharing platforms [17].

Second, a series of government and industry-supported programs on AI4SI has helped the growth of the area. For example, the National Science Foundation in the United States sponsored a research network on computational sustainability [3], which supports researchers from more than 10 universities to develop cross-cutting computational methods that can be applied to sustainability challenges including conservation, poverty mitigation, and renewable energy. Google launched the Google AI Impact Challenge in 2018 [4] and then launched AI for Social Good Workshop in 2021 and 2022 [1] which provides funding for researchers and their partners from non-government organizations (NGOs) to work on AI4SI projects. The XPrize Foundation launched an initiative on AI and Data for Good [2] to support and inspire innovative and scalable work to solve societal grand challenges.

Third, there have been a few workshops, conferences, and special tracks that specifically focused on AI4SI, which raised awareness and attracted the attention of the research community on AI4SI. For example, in 2016, the Artificial Intelligence For Social Good workshop was held with support from the Association for the Advancement of Artificial Intelligence (AAAI), the Computing Community Consortium (CCC), and the White House Office of Science Technology Policy (OSTP). This workshop seemed to be a key signal that triggered other workshops and symposiums on AI4SI have been held in AI conferences in the past few years, e.g., the AAAI 2017 Spring Symposium on AI for Social Good, and the AI for Social Good workshop at NeurIPS 2019. More recently, AAAI and IJCAI, two top-tier AI conferences, started their special tracks on AI4SI and AI for Social Good in 2020 and 2021 respectively.

A survey paper [20] investigated more than 1000 conference papers from major AI conferences relevant to AI4SI during 2008-2019. Not surprisingly, the number of papers increased from 18 in 2008 to 246 in 2019. These papers develop AI methods to tackle challenges in many different domains, including healthcare, transportation, agriculture, environmental sustainability, education, social care, urban planning, and many others. As observed in the survey paper, healthcare has received the most attention among all application domains. The difference between it and other domains appears to be widening, with the works on healthcare counting 32% of the total relevant papers in 2019. Transportation ranks second among all the domains in the total number of papers in the past years. It is not surprising as many problems in healthcare and transportation can be formulated as well-defined optimization problems and there has been a lot of

work on developing computational methods for solving these optimization problems without using AI techniques. For example, the traveling salesman problem (TSP) [12] has been a long-lasting research topic in operations research, and has many direct applications in the transportation domain. The development of AI inspires many new algorithms for TSP that are more scalable [23, 18], which can help solve important tasks in transportation such as travel itinerary planning. Also, many datasets are readily available in healthcare and transportation or can be collected relatively readily in collaboration with stakeholders thanks to existing digital infrastructure such as electronic health records and traffic monitoring systems. In contrast, agriculture, education, social care, and urban planning have been the topics with the fewest papers. Data in these domains are often in forms like interviews and questionnaires, which make it much harder to collect and parse the data.

In terms of the AI techniques used in these AI4SI-related conference papers, machine learning (ML) has been a consistently dominant technique, and its usage has surpassed that of other techniques even more since 2013 as reported in the survey paper [20]. CV and NLP come second, which is not surprising given the breakthroughs in these two areas in the deep learning era and their successful commercial use for many problems. Human-AI collaboration has received more and more attention recently, as ensuring a significant positive social impact often requires AI and human beings to work together in an efficient and meaningful way.

4 Structure of the Book

This book has three parts. Part I is this introductory chapter which provides an overview of AI4SI. We hope that this introductory chapter explains what is AI4SI, why AI4SI differs from traditional AI, and how the research efforts on AI4SI emerge and grow over time.

Part II presents a series of case studies on AI4SI. These case studies serve as examples of how AI has been used to help address societal challenges and how AI can make a significant social impact in the real world in various domains. All of these case studies have led to field tests and some of them have been successfully sustainably deployed in the field. The case studies are grouped based on their application domain. For this version of the book, we focus on four domains: (i) public health; (ii) agriculture and food security; (iii) environment and conservation; and (iv) crisis management and disaster response. We have collected 2–4 case studies for each domain. It is important to note that the success of many of these projects is due to joint efforts between AI researchers and practitioners who have been working in these domains for years.

Each case study chapter is structured to include the following sections:

- Introduction & Problem Statement: These sections cover a description of the target problem, the previous practice before AI is used, and why is AI needed for the problem.

- Method: This section covers the AI frameworks, models, algorithms, and other important parts in the development of the project, including new methods for data collection.
- Resource Requirements: This section covers datasets used, computing resources needed, and other tools (open sourced or proprietary) used in the development, deployment, and maintenance of the project.
- Field Evaluation: This section covers the deployment status of the project and the results.
- Lessons learned: This section covers lessons learned in the development, deployment, and maintenance of the project and can include both positive and negative aspects.

Part III of the book discusses fundamental challenges in AI4SI across many problems and domains and lessons learned from research fields related to AI4SI. Specifically, we will have chapters discussing how to properly design experiments and evaluate AI4SI work and how to ensure a user-centered design to develop AI4SI projects. We will introduce the lessons from HCI, ICTD, and their implications for AI4SI.

We expect more case studies and discussions on the fundamental challenges of AI4SI in a future version of the book.

5 Scope and Target Audience

This book provides an overview of AI4SI, introduces a sample of existing efforts on AI4SI through case studies, and presents several fundamental challenges in AI4SI and lessons from related fields. We are aware of many other existing efforts and great work on AI4SI. For readers who would like to learn more about other existing work, we refer the readers to the additional resources page of the book¹. We also expect future versions of the book to have enriched content.

We attempt to present the book in a fashion that is accessible to students and researchers who have some preliminary understanding of the basics of AI. We also believe that the case studies in Part II of the book would be of interest to students and researchers who have worked on the corresponding domains ranging from public health to disaster response. We expect the book to serve as a textbook for classes on AI4SI.

References

1. 2022 ai for social good workshop, google. <https://sites.google.com/view/ai4sgprogram2022>, accessed: 2022-08-25
2. Ai and data for good, xprize. <https://www.xprize.org/aidataforgood>, accessed: 2022-08-25
3. Computational sustainability network. <http://www.compsust.net/>, accessed: 2022-02-11

¹ <https://ai4sibook.org/reading-materials>

4. Google ai impact challenge, 2018. <https://impactchallenge.withgoogle.com/ai2018>, accessed: 2022-08-25
5. Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al.: Dota 2 with large scale deep reinforcement learning. arXiv preprint arXiv:1912.06680 (2019)
6. Blount, D., Gero, S., Van Oast, J., Parham, J., Kingen, C., Scheiner, B., Stere, T., Fisher, M., Minton, G., Khan, C., et al.: Flukebook: an open-source ai platform for cetacean photo identification. *Mammalian Biology* pp. 1–19 (2022)
7. Brown, N., Sandholm, T.: Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* p. eaao1733 (2017)
8. Coles, L.S.: The application of artificial intelligence to medicine. *Futures* **9**(4), 315–323 (1977)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
11. Dou, Z.Y., Barman-Adhikari, A., Fang, F., Yadav, A.: Harnessing social media to identify homeless youth at-risk of substance use. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 14748–14756 (2021)
12. Jünger, M., Reinelt, G., Rinaldi, G.: The traveling salesman problem. *Handbooks in operations research and management science* **7**, 225–330 (1995)
13. Minsky, M.: Steps toward artificial intelligence. *Proceedings of the IRE* **49**(1), 8–30 (1961)
14. OpenAI: OpenAI Five. <https://blog.openai.com/openai-five/> (2018)
15. Parham, J., Stewart, C., Crall, J., Rubenstein, D., Holmberg, J., Berger-Wolf, T.: An animal detection pipeline for identification. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1075–1083. IEEE (2018)
16. Perrault, A., Fang, F., Sinha, A., Tambe, M.: Ai for social impact: Learning and planning in the data-to-deployment pipeline. *AI Magazine* (2020)
17. Qin, Z., Zhu, H., Ye, J.: Reinforcement learning for ridesharing: An extended survey. arXiv e-prints pp. arXiv–2105 (2021)
18. Rani, S., Kholidah, K.N., Huda, S.N.: A development of travel itinerary planning application using traveling salesman problem and k-means clustering approach. In: Proceedings of the 2018 7th International Conference on Software and Computer Applications. pp. 327–331 (2018)
19. Schneider, S., Taylor, G.W., Linquist, S., Kremer, S.C.: Past, present and future approaches using computer vision for animal re-identification from camera trap data. *Methods in Ecology and Evolution* **10**(4), 461–470 (2019)
20. Shi, Z.R., Wang, C., Fang, F.: Artificial intelligence for social good: A survey. arXiv preprint arXiv:2001.01818 (2020)
21. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of Go with deep neural networks and tree search. *Nature* **529**(7587), 484–489 (2016)
22. Vinyals, O., Babuschkin, I., Czarnecki, W.M., Mathieu, M., Dudzik, A., Chung, J., Choi, D.H., Powell, R., Ewalds, T., Georgiev, P., et al.: Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature* **575**(7782), 350–354 (2019)

23. Xin, L., Song, W., Cao, Z., Zhang, J.: Neurolkh: Combining deep learning model with lin-kernighan-helsgaun heuristic for solving the traveling salesman problem. *Advances in Neural Information Processing Systems* **34**, 7472–7483 (2021)