

Lessons from the 2020 AAAI Fall Symposium on AI for Social Good

Oshani Seneviratne¹, Hemant Purohit², Muhammad Aurangzeb Ahmad³

¹Rensselaer Polytechnic Institute, USA

²George Mason University, USA

³University of Washington Tacoma and KenSci Inc., USA

¹senevo@rpi.edu, ²hpurohit@gmu.edu, ³maahmad@uw.edu

Abstract

Recent developments in big data and computational power are revolutionizing several domains, opening up new opportunities and challenges. In this symposium, we highlighted two specific themes, namely humanitarian relief and healthcare, where AI could be used for social good to achieve the United Nations (UN) sustainable development goals (SDGs) in those areas, which touch every aspect of human, social, and economic development. The talks at the symposium were focused on identifying the critical needs and pathways for responsible AI solutions to achieve SDGs, which demand holistic thinking on optimizing the trade-off between automation benefits and their potential side-effects, especially in a year that has upended societies globally due to the COVID-19 pandemic.

Introduction

Riding on the success of the AI for Social Good symposium that was held in Washington, DC, in November 2019, we organized the 2020 version of the symposium. While keeping the focus on the two UN SDGs of healthcare and disaster relief, we strove to highlight the challenges of trust deficit as the cost of AI errors and how to do responsible AI system design as part of the 2020 symposium. We identified several directions of AI for social good, including the reliability and robustness guarantees, human-centered approach for testing, ethical design, explainability, fairness, and the elimination of AI bias.

Given the unique circumstances in 2020, the symposium's two themes were uniquely fitting for this year. The worldwide healthcare crisis with the COVID-19 pandemic was akin to a *Grey Rhino* event (i.e., highly probable but neglected threat that has an enormous impact). Simultaneously, various disaster scenarios such as wildfires early in the year in Australia and later in the year in the Western United States requiring humanitarian relief characterized the *Black Swan* event (i.e., an unpredictable event beyond what is typically expected of a situation and has potentially severe consequences).

AAAI Fall 2020 Symposium on AI for Social Good.

Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Objectives

The key objectives of the symposium along the two key themes are as follows.

AI for Healthcare: Healthcare is one of the foremost challenges of today's world, highlighted by the recent COVID-19 pandemic where it has come to the forefront of the global discourse. In general, healthcare data is characterized by data missingness, poor data standardization, data incompleteness, and other data quality issues that have downstream consequences. These factors hinder the deployment of solutions relevant to real-world use cases. Moreover, AI, particularly Machine Learning (ML), system design in healthcare is characterized by the last mile problem, where delivering a practical solution that is reliable and robust to errors (especially in "break glass in case of emergency" situations) has proven hard to implement. These have broader implications in the context of fairness, explainability, and transparency in ML. Therefore, the implementation and deployment of AI/ML systems in healthcare bring up challenges that go far beyond model building and scoring. This symposium also focused on a broad range of AI healthcare applications and challenges encountered, including but not limited to: automation bias, prescriptive AI models, explainability, privacy and security, transparency, and decision rights, especially in the context of deployment of AI in real-world scenarios in healthcare.

AI for Humanitarian Technologies and Disaster Management: Technology can have an incredible impact on how we address humanitarian issues and achieve SDGs worldwide. Detecting and predicting how a crisis or conflict could develop, analyzing the impact of catastrophes in a cyber-physical society, and assisting in disaster response and resource allocation are of utmost importance, where the advances in AI can be utilized. The AI techniques can allow better preparation for disasters, help save lives, limit economic losses, provide adequate disaster relief, and make communities more robust and resilient. The symposium focused on all aspects of humanitarian relief operations supported by the novel use of AI technologies from enabling missing persons to be located, leveraging crowdsourced data to provide early warning for rapid response to emergencies, increasing situational awareness, to logistics and supply chain management.

Program

Keynotes

We had several esteemed researchers and thought leaders deliver several keynotes at the workshop. Our opening keynote speaker was Prof. Malik Magdon-Ismael from the Rensselaer Polytechnic Institute, who gave an insightful keynote on simple local models with robust change-point analysis and model identification for COVID-19 prediction that can be applied at the county or organization level. Highlighting the importance of AI explainability, Dr. Rich Caruana from Microsoft Research talked about glass box models in an aptly titled keynote “Friends Don’t Let Friends Deploy Black-Box Models: The Importance of Intelligibility in Machine Learning.” We were fortunate to have Hon. Maleeh Jamal, the Minister for Communication, Science, and Technology of The Maldives, deliver the opening keynote of the second day on “AI for Social Good: Small Nations’ Perspectives.” Dr. Walter Dorn from the Royal Military College of Canada & Canadian Forces College presented a vision on “Intelligence for Peace: AI in UN Field Operations and Cyber-peacekeeping,” providing perspectives from various exciting use cases. Our closing keynote was by Dr. Suranga Nanayakkara from the University of Auckland, New Zealand, who talked about inspiring tools and techniques to augment human capabilities, focusing on assistive technologies.

Panels

The symposium had two panels, one focused on AI for healthcare and the other on AI for Humanitarian Technologies and Disaster Management.

The panel on AI for Healthcare, moderated by Dr. Muhammad Aurangzeb Ahmad consisted of Dr. Carly Eckert, MD (Department of Epidemiology, University of Washington & KenSci), Dr. Vikas Kumar (KenSci), Dr. Nicholas Mark, MD (Swedish Hospital), and Dr. Oshani Seneviratne (Rensselaer Polytechnic Institute). The panelists discussed the last mile problem in healthcare AI, which is the challenge of adopting and implementing ML models in the clinical workflow, highlighting the main hurdles that need to be overcome in their opinion. The conversation also focused on what roles, if any, AI can play in reducing delivery bias because while data bias and algorithmic bias are relatively straightforward to quantify, the delivery bias in healthcare is trickier to quantify. The discussion also included what regulatory bodies should focus on given the rapid pace of technological progress in AI/ML, and more importantly, whether such technologies can be regulated meaningfully. Especially given that underserved and underprivileged communities often do not have access to the tools even to know if they are being discriminated against, the panel discussed what the AI community, with cooperation from the healthcare practitioners, can do to remedy these problems along with many insights from their work on applying AI in healthcare in real-world settings.

The panel on AI for Humanitarian Technologies and Disaster Management, moderated by Dr. Hemant Purohit, consisted of Dr. Jennifer Chan, MD, MPH (Professor, Feinberg

School of Medicine, Northwestern University), Mr. Steve Peterson, CEM (Montgomery County CERT, and National Institutes of Health), Dr. Walter Dorn (Professor, Royal Military College of Canada and United Nations Peacekeeping Operations), and Dr. Oshani Seneviratne (Rensselaer Polytechnic Institute). The panelists shared success stories of using AI technology in humanitarian assistance and disaster management. The discussion then shifted to potential barriers for adopting AI in this space, both in terms of operational and data or technology-centric challenges. The panelists identified the concerns of limited capabilities and the need for AI tools to reach and respond to the last mile during disaster relief, such as diverse speaking populations and remote vulnerable areas with conflicts.

Paper Selection

We received 28 papers from 68 authors to the call for papers. After a rigorous peer-review process with the help of our program committee members that consisted of 27 researchers from a variety of research areas, we selected 22 papers as regular papers and 3 papers as short papers. Each paper received at least two reviews. In terms of the topics, we had a variety of novel research spanning healthcare and humanitarian technologies. Unsurprisingly, we had many papers on COVID-19 related topics ranging from applications of policy guidance and mitigation from epidemiological data to using computer vision to ascertain that individuals follow social distancing guidelines. The first day of the symposium was dedicated to discussing healthcare technologies, and the second day for humanitarian technologies.

Audience Participation

Unlike the previous year, due to the COVID-19 pandemic, we decided to hold the symposium virtually. We used this as an opportunity to increase participation, as those who would not usually be able to travel to Washington, DC., would now be able to attend the symposium. We saw participants from all over the USA, as well as from around the world.

Website

The symposium details, including the program, keynote speakers and the panelists, and the recorded videos of all the paper presentations are available on our website at <https://ai-for-socialgood.github.io/2020/index.html>.

Conclusion

The AI for Social Good Fall-2020 symposium was built upon our continued efforts in bringing the AI community members together for the healthcare and humanitarian technology themes and reinforced the success of last year’s successful AAIAI Fall Series symposium on AI for Social Good. The symposium brought together AI researchers, domain scientists, practitioners, and policymakers to exchange problems and solutions, identify synergies across different application domains, and lead to future collaborative efforts. We will continue to organize similar events to have further discourse on this important topic of AI Social Good.

Organizers

This symposium was co-organized by Dr. Muhammad Aurangzeb Ahmad, Dr. Hemant Purohit, and Dr. Oshani Seneviratne. Dr. Muhammad Aurangzeb Ahmad is an Affiliate Assistant Professor at the University of Washington Tacoma and Principal Research Scientist at KenSci Inc. Dr. Hemant Purohit is an Assistant Professor of Information Sciences and Technology at George Mason University. Dr. Oshani Seneviratne is the Director of Health Data Research at Rensselaer Polytechnic Institute.