Socially Responsible AI Algorithms: Issues, Purposes, and Challenges

Abstract

AI has the potential to drive us towards a future in which all of humanity flourishes. It also comes with substantial risks for oppression and calamity. Technologists and AI researchers have a responsibility to develop trustworthy AI systems. They have responded with great effort to design more responsible AI algorithms. However, existing technical solutions are narrow in scope and have been primarily directed towards algorithms for scoring or classification tasks, with an emphasis on fairness and unwanted bias. To build long-lasting trust between AI and human beings, we argue that the key is to think beyond algorithmic fairness and connect major aspects of AI that potentially cause AI's indifferent behavior. In this tutorial, we examine the subjects and causes of AI indifference, define the objectives of socially responsible AI algorithms, and introduce the means by which we may achieve these objectives. We further discuss how socially responsible AI algorithms improve societal well-being through protection, information, and prevention/mitigation.

Intended Audience

This tutorial is intended for researchers and practitioners who are interested in socially responsible AI and have basic knowledge of AI, data mining, and machine learning. It will be delivered at a college junior/senior level and should be easily accessible to interested parties from both industry and academia.

Tutorial Presenters Bios

1. *Lu Cheng* is a PhD student in Computer Science at ASU. Her research focuses on understanding the relationship between AI and Human from both *statistical* and *causal* perspectives. Lu has published in WSDM, AAAI, IJCAI, and CIKM, among others. Lu won the 2020 ASU Graduate Outstanding Research Award, IBM Ph.D. Social Good Fellowship, Visa Research Scholarship, ASU Grace Hopper Celebration Scholarship, ASU GPSA travel support, and multiple conferences' student scholarships and travel awards. Contact her at lcheng@asu.edu.

2. *Fred Morstatter* is a Research Assistant Professor at USC and a Research Lead at USC's Information Sciences Institute. His research focuses on understanding biases that occur in online social data, forecasting, and identifying cultural models. He is also interested in characterizing the biases of cultural groups based upon the trace data they create on social media. He has published in JMLR, WWW, KDD, Nature, and ICWSM, among others. He was a Program Committee Chair for ICWSM 2019 and Presentation chair for WebSci 2020 and 2021. More information can be found at isi.edu/~fredmors.

3. *Huan Liu* is a professor of Computer Science and Engineering at ASU. His research interests are in data mining, machine learning, social computing, and artificial intelligence. He is a co-author of a text, *Social Media Mining: An Introduction,* Cambridge University Press, Field Chief Editor of Frontiers in Big Data and its Specialty Chief Editor of Data Mining and Management. He is a Fellow of ACM, AAAI, AAAS, and IEEE.

Tutorial Structure

1. Introduction (20 min, Huan Liu)
    - What is socially responsible AI? How does it influence society and people's lives?
    - How is socially responsible AI different from similar concepts such as trustworthy AI and ethical AI?
2. Social Responsibility of AI and the Pyramid (20 min, Lu Cheng)
3. 10-minute break
4. Socially Responsible AI Algorithms (50 min, Fred Morstatter and Lu Cheng)
    - Causes and subjects of socially indifferent AI algorithms
    - Objectives of socially responsible AI algorithms
    - Means
5. 10-minute break
6. Roles of Socially Responsible AI Algorithms (50 min, Fred Morstatter and Lu Cheng)
    - Informing
    - Protecting
    - Preventing
7. Open problems and Challenges (20 min, Huan Liu)