AI Policy

Larry Medsker (George Washington University; lrm@gwu.edu)
DOI: 10.1145/3175502.3175510

Abstract

AI Policy is a regular column in AI Matters featuring summaries and commentary based on postings that appear twice a month in the AI Matters blog (https://sigai.acm.org/aimatters/blog/). Selected posts are summarized in issues of *AI Matters*.

Introduction

The SIGAI Public Policy goals are to:

- promote discussion of policies related to AI through posts in the AI Matters blog on the 1st and 15th of each month,

- help identify external groups with common interests in AI Public Policy,

- encourage SIGAI members to partner in policy initiatives with these organizations, and

- disseminate public policy ideas to the SIGAI membership through articles in the newsletter.

I welcome everyone to make blog comments so we can develop a rich knowledge base of information and ideas representing the SIGAI members.

Organizations Related to AI and Policy

In 2017 we expanded our reporting on and work with the American Association for the Advancement of Science, particularly the Center of Science, Policy, and Society. While AAAS policy issues are usually not directly related to AI, a regular look at their Policy Alert notifications is useful for larger policy issues, and helpful to see opportunities for SIGAI to be involved in public policy events. We also expanded our relationship with the ACM US Public Policy Council (USACM), which serves as the focal point for ACM's interaction with US government organizations, the computing

community, and the US public in all matters of US public policy related to information technology. USACM addresses issues in innovation, privacy, security, digital governance, intellectual property, accessibility, and e-voting. I am a member of USACM, in part representing our ACM SIGAI.

Algorithmic Transparency and Accountability

ATA is a current major initiative with USACM. Your Public Policy Officer attended and reported on the USACM Panel on Algorithmic Transparency and Accountability, ehcih took place on Thursday, September 14th at the National Press Club. The panelists were moderator Simson Garfinkel, Jeanna Neefe Matthews, Nicholas Diakopoulos, Dan Rubins, Geoff Cohen, and Ansgar Koene. USACM Chair Stuart Shapiro opened the event, and Ben Sneiderman provided comments from the audience.

USACM and EUACM have identified and codified a set of principles intended to ensure fairness in this evolving policy and technology ecosystem. These were a focus of the panel discussion and are as follows: (1) awareness; (2) access and redress; (3) accountability; (4) explanation; (5) data provenance; (6) auditability; and (7) validation and testing. See also the full letter in the September 2017 issue of CACM.

The panel and audience discussion ranged from frameworks for evaluating algorithms and creating policy for fairness to examples of algorithmic abuse. Language for clear communication with the public and policymakers, as well as even scientists, was a concern—as has been discussed in our Public Policy blog. Algorithms in the strict sense may not always be the issue, but rather the data used to build and train a system, especially when the system is used for prediction and decision making. Much was said about the types of bias and unfairness that can be embedded in modern AI and machine learning systems.

The essence of the concerns includes the ability to explain how a system works, the need to develop models of algorithmic transparency, and how policy or an independent clearinghouse might identify fair and problematic algorithmic systems. Please read more about the panel discussion at https://www.acm.org/public-policy/algorithmic-panel and view the video at https://www.youtube.com/watch?v=DDW-nM8idgg&feature=youtu.be.

AAAI Fall 2017 Symposium Series

This year's Fall Symposium Series (November 9-11) https://aaai.org/Symposia/Fall/fss17symposia.php provided updates and insights on advances in research and technology, including resources for discussion of AI policy issues. The symposia addressed topics in human-robot interaction, cognitive assistance in government and public sectors, military applications, human-robot collaboration, and a standard model of the mind. Important themes for public policy were about the advances and questions on human-AI collaboration.

The cognitive assistance sessions this year focused on government and public sector applications, particularly autonomous systems, healthcare, and education. Human-technology collaboration advances involved discussions of issues relevant to public policy, including privacy and algorithmic transparency. The increasing mix of AI with humans in ubiquitous public and private systems was the subject of discussions about new technological developments and the need for understanding and anticipating challenges for communication and collaboration. Particular issues were on jobs and de-skilling of the workforce, credit and blame when AI applications work or fail, and the role of humans with autonomous systems.

IBM's Jim Spohrer made an outstanding presentation "A Look Toward the Future," incorporating his rich experience and current work on anticipated impacts of new technology. His slides are well worth studying, especially for the role of hardware in game-changing technologies with likely milestones every ten years through 2045. Radical developments in technology would challenge public policy in ways that are difficult to imagine, but current poli-

cymakers and the AI community need to try. See related references in the Resources section below.

Particular takeaways, and anticipated subjects for future blogs, are about the importance of likely far-reaching research and applications on public policy. The degree and nature of cognitive collaboration with machines, the future of jobs, new demands on educational systems as cognitive assistance becomes deep and pervasive, and the anticipated radical changes in AI capabilities put the challenges to public policy in a new perspective. AI researchers and developers need to partner with social scientists to anticipate communication and societal issues as human-machine collaboration accelerates, both in system development teams and in the new workforce.

Collaborations with other Policy and Ethics Groups

SIGAI recently began discussion with other groups, particularly between ACM and IEEE, on finding ways to bring together efforts on Algorithmic Transparency and Accountability. One opportunity is at RightsCon Toronto: May 16-18, 2018. The call for proposals mentioned "Artificial Intelligence, Automation, and Algorithmic Accountability' as one of their program "buckets."

As AI is becoming more pervasive in our lives, the impact on society is increasingly significant. Concerns and issues are being raised regarding value alignment, data bias and data policy, regulations, and workforce displacement. We need multi-disciplinary and multi-stakeholder efforts to find the best ways to address concerns using expertise from AI, computer science, ethics, philosophy, economics, sociology, psychology, law, history, and politics. AAAI and ACM are joining forces to start a new conference, the AAAI/ACM Conference on AI, Ethics, and Society. The first edition of this conference http://www.aies-conference.com will be co-located with AAAI-18 on February 2-3, 2018, in New Orleans. The program of the conference will include peer-reviewed paper presentations, invited talks, panels, and working sessions. The conference will cover a broad set of topics, including trust and explanations in AI systems, fairness and transparency, ethical design and development of AI

systems, and impact of AI on the workforce.

Upcoming

Some themes planned for the SIGAI Public Policy posts for 2018 include algorithmic accountability, human-machine collaboration, and the impacts of AI and Data Science on the future of education and the labor market. We will look at potential policies for today that could mitigate impacts of AI on individuals and society. Policy areas include innovative educational systems, ideas for alternate economic systems, and regulatory changes to promote safe and fair technological innovation.

Resources

- AAAI information related to science policy issues:
  https://aitopics.org/search

- AAAI Symposium Series:
  https://aaai.org/Symposia/Fall/fss17.php

- Ansgar Koene:
  https://theconversation.com/
  machine-gaydar-ai-is-reinforcing-stereotypes-that-liberal-societies-are-trying-to-get-rid-of-83837

- CACM letter on algorithmic transparency and accountability:
  https://cacm.acm.org/magazines/2017/9/
  220423-toward-algorithmic-transparency-and-accountability/
  fulltext#FNA

- Jim Sphorer, A Look Toward the Future:
  https://www.slideshare.net/spohrer

- Humans, robotics, and the future of manufacturing:
  https://www.engadget.com/2017/09/11/
  human-robot-ai-manufacturing/

- New education systems and the future of work:
  https://www.edweek.org/ew/articles/2017/09/
  27/the-future-of-work-is-uncertain-schools.
  html

- Noriko Arai's TED talk on Todai Robot:
  https://www.ted.com/talks/noriko_arai_can_a_
  robot_pass_a_university_entrance_exam

- Smart phone app "Seeing AI":
  https://www.youtube.com/watch?v=
  bqeQByqf_f8

Larry Medsker is Research Professor of Physics and Director of the Data Science graduate program at The George Washington University. Dr. Medsker is a former Dean of the Siena College School of Science, and Professor in Computer Science and in Physics, where he was a co-founder of the Siena Institute for Artificial Intelligence. His research and teaching continues at GW on the nature of humans and machines and the impacts of AI on society and policy[a]. Professor Medskers research in AI includes work on artificial neural networks and hybrid intelligent systems. He is the Public Policy Officer for the ACM SIGAI.

[a]http://www.humai.org/humai/ and http://humac-web.org/