Toward Algorithmic Transparency and Accountability

ALGORITHMS ARE REPLACING or augmenting human decision making in crucial ways. People have become accustomed to algorithms making all manner of recommendations, from products to buy, to songs to listen to, to social network connections. However, algorithms are not just recommending, they are also being used to make big decisions about people's lives, such as who gets loans, whose résumés are reviewed by humans for possible employment, and the length of prison terms. While algorithmic decision making can offer benefits in terms of speed, efficiency, and even fairness, there is a common misconception that algorithms automatically result in unbiased decisions. In reality, inscrutable algorithms can also unfairly limit opportunities, restrict services, and even improperly curtail liberty.

Information and communication technologies invariably raise these kinds of important public policy issues. How should self-driving cars be required to act? How private is information stored on a cellphone? Can electronic voting machines be trusted? How will the increasing uses of automation in the workplace impact workers? Since its founding, ACM's members have played a leading role in discussing these issues within the computing profession and with policymakers.

The ACM U.S. Public Policy Council (USACM) was established in the early 1990s as a focal point for ACM's interactions with U.S. government organizations, the computing community, and the public in all matters of U.S. public policy related to information technology. USACM came to prominence during the debates over cryptography and key escrow technology. Today, USACM continues to make public policy recommendations that are based on scientific evidence, follow recognized best practices in computing, and are grounded in the ACM Code of Ethics. It has established a reputation as a non-partisan,

principled, and independent source of scientific and technical expertise, free from the influence of product vendors or other vested interests.

More recently, the ACM Europe Council Policy Committee (EUACM) has been doing the same in Europe. USACM and EUACM, both separately and jointly, provide information and analysis to policymakers and the public regarding important societal issues involving IT, including *algorithmic transparency and accountability*.

USACM and EUACM have identified and codified a set of principles intended to ensure fairness in this evolving policy and technology ecosystem.^a These are: (1) awareness; (2) access and redress; (3) accountability; (4) explanation; (5) data provenance; (6) auditability; and (7) validation and testing.

Awareness speaks to educating the public regarding the degree to which decision making is automated. *Access and redress* means there is a way to investigate and correct erroneous decisions. *Accountability* rejects the common deflection of blame to an automated system by ensuring those who deploy an algorithm cannot eschew responsibility for its actions. *Explanation* means the logic of the algorithm, no matter how complex, must be communicable in human terms.

As many modern techniques are based on statistical analyses of large pools of collected data, decisions will be influenced by the choice of datasets for training, and thus knowing the data sources and their trustworthiness—that is, their *provenance*—is essential. *Auditability* for a decision-making system requires logging and record keeping, for example, for dispute resolution or regulatory compliance. Finally, *validation and testing* on an ongoing basis means that techniques such as regression tests, vetting of corner cases, or red-teaming strate-

gies used in computer security should be employed to increase confidence in automated systems.

As organizations deploy complex algorithms for automated decision making, system designers should build these principles into their systems. In some cases, doing so will require additional research. For example, how to design and deploy large-scale neural networks while ensuring compliance with laws prohibiting discrimination against legally protected groups? This is especially crucial given the ability to infer characteristics such as gender, race, or disability status even if the computer system is not provided with that data directly. How should information on automated decisions be logged to ensure auditability? How can the operation of these networks be explained to technologists and non-technical policymakers alike?

One model for moving forward may be self-regulation by industry. Our experience, however, is that self-regulation is only possible when there is a consensus on a set of relevant standards. We hope our principles can serve as input to such an effort. If policymakers determine regulation is necessary, our principles are available, potentially in the way that the Code of Fair Information Practices provided a basis for decades of privacy regulation around the world.

USACM and EUACM seek input and involvement from ACM's members in providing technical expertise to decision makers on the often difficult policy questions relating to algorithmic transparency and accountability, as well as those relating to security, privacy, accessibility, intellectual property, big data, voting, and other technical areas. For more information, visit www.acm.org/public-policy/usacm or www.acm.org/euacm.

The authors are members of the ACM U.S. Public Policy Council, for which **Stuart S. Shapiro** (s_shapiro@acm.org) serves as chair.

Copyright held by authors.

^a https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf