August 19, 2021

COMMENTS IN RESPONSE TO RFI ON NIST ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (DOCKET NUMBER 210726–0151)

The Association for Computing Machinery (ACM), with more than 50,000 U.S. members and approximately 100,000 worldwide, is the world's largest educational and scientific computing society. ACM's US Technology Policy Committee (USTPC), currently comprising more than 130 members, serves as the focal point for ACM's interaction with all branches of the US government, the computing community, and the public on policy matters related to information technology. USTPC is pleased to contribute to the above-referenced proceeding and, as an initial matter, fully agrees with the manner in which it has been framed:

> *[N]ew AI-based technologies, products, and services bring technical and societal challenges and risks, including ensuring that AI comports with ethical values. While there is no objective standard for ethical values, as they are grounded in the norms and legal expectations of specific societies or cultures, it is widely agreed that AI must be designed, developed, used, and evaluated in a trustworthy and responsible manner to foster public confidence and trust. Trust is established by ensuring that AI systems are cognizant of and are built to align with core values in society, and in ways which minimize harms to individuals, groups, communities, and societies at large.*

To these ends, USTPC urges NIST to consider, as appropriate, ACM's benchmark Code of Ethics and Professional Conduct [https://ethics.acm.org]. Two years in the making, the Code was revised in 2018 with extensive input from ACM members worldwide. We also respectfully commend to NIST's attention the attached Statement on Algorithmic Transparency and Accountability. Endorsed jointly by both ACM's Europe and U.S. policy committees, the Statement articulates seven guiding principles intended to inform key aspects of algorithm design and deployment.

USTPC believes that both these documents can be of foundational benefit to the effort articulated in this proceeding and stands ready to objectively and apolitically facilitate NIST's deliberations and analyses. To arrange for technical and apolitical input from USTPC and/or ACM's expert members, please contact Adam Eisgrau, ACM Director of Global Policy & Public Affairs, at acmpo@acm.org or 202-580-6555.

ACM U.S. Technology Policy Committee
Pennsylvania Ave NW, Suite 200
Washington, DC 20006

+1 202.580.6555 1701
acmpo@acm.org
www.acm.org/public-policy/ustpc

Updated May 25, 2017

Statement on Algorithmic Transparency and Accountability

by ACM U.S. Public Policy Council, approved January 12, 2017
ACM Europe Policy Committee, approved May 25, 2017

Computer algorithms are widely employed throughout our economy and society to make decisions that have far-reaching impacts, including their applications for education, access to credit, healthcare, and employment. The ubiquity of algorithms in our everyday lives is an important reason to focus on addressing challenges associated with the design and technical aspects of algorithms and preventing bias from the onset.

An algorithm is a self-contained step-by-step set of operations that computers and other 'smart' devices carry out to perform calculation, data processing, and automated reasoning tasks. Increasingly, algorithms implement institutional decision-making based on analytics, which involves the discovery, interpretation, and communication of meaningful patterns in data. Especially valuable in areas rich with recorded information, analytics relies on the simultaneous application of statistics, computer programming, and operations research to quantify performance.

There is also growing evidence that some algorithms and analytics can be opaque, making it impossible to determine when their outputs may be biased or erroneous.

Computational models can be distorted as a result of biases contained in their input data and/or their algorithms. Decisions made by predictive algorithms can be opaque because of many factors, including technical (the algorithm may not lend itself to easy explanation), economic (the cost of providing transparency may be excessive, including the compromise of trade secrets), and social (revealing input may violate privacy expectations). Even well-engineered computer systems can result in unexplained outcomes or errors, either because they contain bugs or because the conditions of their use changes, invalidating assumptions on which the original analytics were based.

The use of algorithms for automated decision-making about individuals can result in harmful discrimination. Policymakers should hold institutions using analytics to the same standards as institutions where humans have traditionally made decisions and developers should plan and architect analytical systems to adhere to those standards when algorithms are used to make automated decisions or as input to decisions made by people.

This set of principles, consistent with the ACM Code of Ethics, is intended to support the benefits of algorithmic decision-making while addressing these concerns. These principles should be addressed during every phase of system development and deployment to the extent necessary to minimize potential harms while realizing the benefits of algorithmic decision-making.

ACM U.S. Technology Policy Committee
Pennsylvania Ave NW, Suite 200
Washington, DC 20006

+1 202.580.6555 1701
acmpo@acm.org
www.acm.org/public-policy/ustpc

Principles for Algorithmic Transparency and Accountability

1. Awareness: Owners, designers, builders, users, and other stakeholders of analytic systems should be aware of the possible biases involved in their design, implementation, and use and the potential harm that biases can cause to individuals and society.

2. Access and redress: Regulators should encourage the adoption of mechanisms that enable questioning and redress for individuals and groups that are adversely affected by algorithmically informed decisions.

3. Accountability: Institutions should be held responsible for decisions made by the algorithms that they use, even if it is not feasible to explain in detail how the algorithms produce their results.

4. Explanation: Systems and institutions that use algorithmic decision-making are encouraged to produce explanations regarding both the procedures followed by the algorithm and the specific decisions that are made. This is particularly important in public policy contexts.

5. Data Provenance: A description of the way in which the training data was collected should be maintained by the builders of the algorithms, accompanied by an exploration of the potential biases induced by the human or algorithmic data-gathering process. Public scrutiny of the data provides maximum opportunity for corrections. However, concerns over privacy, protecting trade secrets, or revelation of analytics that might allow malicious actors to game the system can justify restricting access to qualified and authorized individuals.

6. Auditability: Models, algorithms, data, and decisions should be recorded so that they can be audited in cases where harm is suspected.

7. Validation and Testing: Institutions should use rigorous methods to validate their models and document those methods and results. In particular, they should routinely perform tests to assess and determine whether the model generates discriminatory harm. Institutions are encouraged to make the results of such tests public.

ACM U.S. Technology Policy Committee
Pennsylvania Ave NW, Suite 200
Washington, DC 20006

+1 202.580.6555 1701
acmpo@acm.org
www.acm.org/public-policy/ustpc