



NEWS RELEASE

Contact: Jim Ormond
ACM Media Relations
212-626-0505
ormond@hq.acm.org

Adam Eisgrau
ACM Director of Global Policy and Public Affairs
202-580-6555
eisgrau@hq.acm.org

World’s Largest Association of Computing Professionals Issues Principles for Generative AI Technologies

Group Claims Existing Mechanisms for Avoiding AI Harms Will Likely Not Suffice

New York, NY, July 11, 2023 – In response to major advances in Generative AI technologies—as well as the significant questions these technologies pose in areas including intellectual property, the future of work, and even human safety—the Association for Computing Machinery’s global Technology Policy Council (ACM TPC) has issued “[Principles for the Development, Deployment, and Use of Generative AI Technologies.](#)”

Drawing on the deep technical expertise of computer scientists in the United States and Europe, the ACM TPC statement outlines eight principles intended to foster fair, accurate, and beneficial decision-making concerning generative and all other AI technologies. Four of the principles are specific to Generative AI, and an additional four principles are adapted from the TPC’s 2022 “[Statement on Principles for Responsible Algorithmic Systems.](#)”

The Introduction to the new Principles advances the core argument that “the increasing power of Generative AI systems, the speed of their evolution, broad application, and potential to cause significant or even catastrophic harm, means that great care must be taken in researching, designing, developing, deploying, and using them. Existing mechanisms and modes for avoiding such harm likely will not suffice.”

The document then sets out these eight instrumental principles, outlined here in abbreviated form:

Generative AI-Specific Principles

- **Limits and guidance on deployment and use:** In consultation with all stakeholders, law and regulation should be reviewed and applied as written or revised to limit the deployment and use of Generative AI technologies when required to minimize harm. No high-risk AI system should

be allowed to operate without clear and adequate safeguards, including a “human in the loop” and clear consensus among relevant stakeholders that the system's benefits will substantially outweigh its potential negative impacts. One approach is to define a hierarchy of risk levels, with unacceptable risk at the highest level and minimal risk at the lowest level.

- **Ownership:** Inherent aspects of how Generative AI systems are structured and function are not yet adequately accounted for in intellectual property (IP) law and regulation.
- **Personal data control:** Generative AI systems should allow a person to opt out of their data being used to train a system or facilitate its generation of information.
- **Correctability:** Providers of Generative AI systems should create and maintain public repositories where errors made by the system can be noted and, optionally, corrections made.

Adapted Prior Principles

- **Transparency:** Any application or system that utilizes Generative AI should conspicuously disclose that it does so to the appropriate stakeholders.
- **Auditability and contestability:** Providers of Generative AI systems should ensure that system models, algorithms, data, and outputs can be recorded where possible (with due consideration to privacy), so that they may be audited and/or contested in appropriate cases.
- **Limiting environmental impact:** Given the large environmental impact of Generative AI models, we recommend that consensus on methodologies be developed to measure, attribute, and actively reduce such impact.
- **Heightened security and privacy:** Generative AI systems are susceptible to a broad range of new security and privacy risks, including new attack vectors and malicious data leaks, among others.

“Our field needs to tread carefully with the development of Generative AI because this is a new paradigm that goes significantly beyond previous AI technology and applications,” explained Ravi Jain, Chair of the ACM Technology Policy Council’s Working Group on Generative AI and lead author of the Principles. “Whether you celebrate Generative AI as a wonderful scientific advancement or fear it, everyone agrees that we need to develop this technology responsibly. In outlining these eight instrumental principles, we’ve tried to consider a wide range of areas where Generative AI might have an impact. These include aspects that have not been covered as much in the media, including environmental considerations and the idea of creating public repositories where errors in a system can be noted and corrected.”

“These are guidelines, but we must also build a community of scientists, policymakers, and industry leaders who will work together in the public interest to understand the limits and risks of Generative AI

as well as its benefits. ACM’s position as the world’s largest association for computing professionals makes us well-suited to foster that consensus and look forward to working with policy makers to craft the regulations by which Generative AI should be developed, deployed, but also controlled,” added James Hendler, Professor at Rensselaer Polytechnic Institute and Chair of ACM’s Technology Policy Council.

“Principles for the Development, Deployment, and Use of Generative AI Technologies” was jointly produced and adopted by ACM’s US Technology Policy Committee (USTPC) and Europe Technology Policy Committee (Europe TPC).

Lead authors of this document for USTPC were Ravi Jain, Jeanna Matthews, and Alejandro Saucedo. Important contributions were made by Harish Arunachalam, Brian Dean, Advait Deshpande, Simson Garfinkel, Andrew Grosso, Jim Hendler, Lorraine Kisselburgh, Srivatsa Kundurthy, Marc Rotenberg, Stuart Shapiro, and Ben Shneiderman. Assistance also was provided by Ricardo Baeza-Yates, Michel Beaudouin-Lafon, Vint Cerf, Charalampos Chelmis, Paul DeMarinis, Nicholas Diakopoulos, Janet Haven, Ravi Iyer, Carlos E. Jimenez-Gomez, Mark Pastin, Neeti Pokhriyal, Jason Schmitt, and Darryl Scriven.

About the ACM Technology Policy Council

[ACM’s global Technology Policy Council](#) sets the agenda for global initiatives to address evolving technology policy issues and coordinates the activities of ACM’s regional technology policy committees in the US and Europe. It serves as the central convening point for ACM’s interactions with government organizations, the computing community, and the public in all matters of public policy related to computing and information technology. The Council’s members are drawn from ACM’s global membership.

About ACM

[ACM, the Association for Computing Machinery](#), is the world’s largest educational and scientific computing society, uniting computing educators, researchers, and professionals to inspire dialogue, share resources, and address the field’s challenges. ACM strengthens the computing profession’s collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

#