



May 30, 2019

Dr. Ann Drobniś, Director
Computing Community Consortium
c/o Computing Research Association
1828 L Street, NW Suite 800
Washington, DC 20036

Re: Comments of the ACM U.S. Technology Policy Committee on CCC
Draft "20-Year Roadmap for Artificial Intelligence Research in the US"

Dear Dr. Drobniś:

On behalf of the U.S. Technology Policy Committee of the Association for Computing Machinery, and in response to the Computing Community Consortium's ("CCC") call for input, I am pleased to summarize here and append in detail the Committee's recommendations* for further iterating CCC's May 2019 Draft "*20-Year Community Roadmap for Artificial Intelligence Research in the US*" ("Roadmap"). We commend and congratulate CCC and the Roadmap's authors for their efforts, and look forward to continuing to collaborate to strengthen and refine this important product.

As you move forward, the Committee respectfully makes the following broad recommendations, which are both elaborated upon and supplemented in the attached Appendix:

1. To enhance the Roadmap's utility and accessibility to all stakeholders, the press, and the general public: shorten the document itself, expand its Executive Summary, and more fully explain what motivated CCC to produce the Roadmap at this time;
2. Include specific short- and long-term goals in the Roadmap and articulate metrics by which progress toward them may best be judged;
3. Address AI in contexts beyond statistical machine learning, including specifically quantum computing, human-computer interaction (HCI), visualization, and health care;
4. Contextualize calls for K-12 AI-based education within current and future STEAM education efforts, and deepen and clarify the Roadmap's discussion of AI's impact on younger students, especially those at levels K-5;

5. Expressly acknowledge that agency and accountability for the development and effects of AI systems reside in the humans behind AI, not the systems themselves, and articulate the need for many forms of robust research oversight, including both academic Institutional Review Boards *and* their industry equivalents;
6. Highlight broadly that research regarding the ethical challenges raised in and by the development and deployment of AI is particularly important, while also specifically noting that select issues merit intensive attention, such as the use of AI in weapons systems and healthcare, algorithmic bias, and the disparate impact of AI's application on various classes of individuals;
7. Strongly articulate the importance of integrating social and ethical agendas into computer science research and development from the outset of all such efforts (particularly those AI-related), rather than addressing such agendas only after product failure;
8. Focus overtly on matters of diversity and inclusion, specifically by deepening discussions of the impact of AI on disadvantaged and marginalized populations, and by expanding authorship and review of the Roadmap to include representatives of both the academic and non-academic research communities; and
9. Include a new "AI and Security" section in the Roadmap.

Thank you again for your work and the opportunity to participate in the development of the Roadmap. Please contact me directly with any questions you may have regarding these comments, or through Adam Eisgrau, ACM's Director of Global Policy & Public Affairs in Washington.

Sincerely,



James A. Hendler, Chair

Appendix

* The Committee is indebted to the members of its ad hoc working group responsible for generating and organizing this submission: coordinator/primary draftsman Simson Garfinkel, principal editor Lorraine Kisselburgh, and contributing members: Andrew Grosso, Larry Medsker, John Murray, Arnon Rosenthal, Ben Shneiderman, Stuart Shapiro, Eugene Spafford and Ricardo Baeza-Yates.

ACM U.S. TECHNOLOGY POLICY COMMITTEE COMMENTS
ON MAY 2019 CCC ROADMAP FOR AI RESEARCH IN THE U.S.

1. ***To enhance the Roadmap’s utility and accessibility to all stakeholders, the press, and the general public: shorten the document itself, expand its Executive Summary, and more fully explain what motivated CCC to produce the Roadmap at this time.***¹

To enhance the document’s accessibility and utility for its principally lay target audiences, we recommend *shortening* the text of the Roadmap to approximately 80 pages. This can be achieved by moving the more detailed elements to a technical appendix and by omitting (or substantially reducing) advice applicable to computer science in general. This category of information includes recommendations about increasing Science, Technology, Engineering, Art and Mathematics (STEAM) training and creating a more diverse workforce, but see our recommendations on inclusion below. Conversely, we also recommend *expanding* the Executive Summary to approximately eight pages to permit covering every significant topic in the Roadmap.

With regard to the Roadmap’s structure, we believe that the document would be stronger if the reasons it was created and deemed to be timely were more fully explained. AI may be the most transformative technology since the invention of computers, but the broader idea of building machines that exhibit some kind of intelligence has been a motivating goal for scientists and engineers for centuries. We suggest detailing why CCC believes that now is the time to create an AI Roadmap, and why its time horizon was set at 20 years rather than a longer or shorter span of time.

In addition, to take full advantage of the Roadmap format and concept, we urge the drafters in the next iteration to weave this valuable collection of ideas and proposals into a coherent narrative organized chronologically to address shorter-, intermediate- and longer-term considerations. Consideration might profitably also be given to explaining the methods that were used to create the Roadmap and selection criteria for topics. This effort could also be aided by referencing other Roadmaps that have successfully organized academic, industrial and government research.²

Finally, we urge the drafters in the next iteration to contextualize the Roadmap’s analysis and recommendations in a discussion of the motivations for producing and objectives of the Roadmap. We also urge reflecting on major AI efforts over the last few decades and how they succeeded or failed. What lessons were learned to inform our current situation? What new issues, technical or societal, have emerged to make this Roadmap necessary?

¹ We recognize that, in order to both shorten the Roadmap itself and add the significant amount of additional discussion we also propose, it may be necessary and useful to create one or more Appendices to the Roadmap.

² For example, it may be productive to discuss the Roadmap for the Japanese Fifth Generation Project. Started in 1982, the Project developed a research Roadmap, chose technologies, and spent \$400 million over 10 years. That project failed. It could be useful to assess and explain why the current Roadmap (which would presumably spend more money over a longer time period, but has similar lofty goals of improving competitiveness while advancing the state of the art) is more likely to succeed.

2. Include specific short- and long-term goals in the Roadmap and articulate metrics against which progress may be judged.

We believe that the Roadmap will be strengthened and more influential on target audiences if it were to complement its long-term recommendations with specific short-term proposals that are likely to be fruitful, sustainable, and to contribute to a range of positive futures. Specifically, such recommendations should be specific so that they can be of maximum use to funding agencies, major government and industry customers, and major research universities. We also believe that the Roadmap will be strengthened by calling for the establishment of metrics by which to measure the objectives it defines, and by providing examples of such benchmarks.

3. Address AI in contexts beyond statistical machine learning, including specifically quantum computing, human-computer interaction, visualization and health care.

We note that the Roadmap sometimes addresses aspects of AI other than statistical machine learning, but the use of statistical AI appears to be its main focus and animating concern. We recommend that the Roadmap more extensively address opportunities for merging traditional AI (such as planning and knowledge representation) with statistical machine learning. We also propose that the Roadmap explicitly address, or treat in greater depth:

- AI and quantum computing, as improvements in quantum computing may be driven by AI and quantum computers are expected to be capable of quickly training neural networks;
- AI and human-computer interaction (“HCI”), given HCI’s long-established centrality to computer science and, more recently, work towards “human-centered artificial intelligence” and new forms of collaborative work;³
- AI-centric visualization, tools for which have been demonstrated to be of great assistance to AI systems developers in understanding their systems, improving performance, and detecting bugs and bias;⁴ and
- AI and its myriad intersections with health care and insurance, which at minimum implicate cybersecurity, privacy, demographics, genomics, gene-editing, and many associated ethical and legal matters.⁵

³ Establishment of Stanford’s Center on this topic is evidence of the breadth and depth of this emerging discipline. See <http://hai.stanford.edu>.

⁴ See work by Fernanda Viegas and the TensorFlow team at Google at <https://accelerate.withgoogle.com/stories/google-fernanda-viegas-designs-human-centered-ai-for-everyone>

⁵ We respectfully suggest that the final Roadmap’s failure to address the role that AI will play in health care would be a significant and glaring omission.

We also believe that the Roadmap would benefit from more discussion of what is now called artificial general intelligence (“AGI”), including the motivations and objectives of its pioneers. Such information, we believe, should serve as useful predicate to:

- acknowledging that there are other possible models for future computing which emphasize human amplification, augmentation, enhancement, and empowerment to produce widely used tools such as search, email and text, navigation, photography, and other ways to build human self-efficacy;⁶
- discussing alternative views of AGI and its desired or probable effects with particular focus on the presentation and integration of cross-disciplinary perceptions of AGI; and
- articulating the need for and desired effects of a larger research agenda geared toward how AI and human intelligence might most productively work together in the future.

Lastly in this context, whether or not AGI is achievable, discussions of such concepts based in scientific research can be motivating to students and the general public. To that end, the Roadmap also might feature a section devoted to involving non-computer scientists in discussion of the promise of AI. Such “outside” participants might constructively include, for example, representatives of the motion picture, video game and virtual reality industries.

4. Contextualize calls for K-12 AI-based education within current and future STEAM education efforts and clarify its discussion of AI’s impact on younger students, especially those at levels K-5.

Subject to substantial further discussion within CCC, we recommend that the call for development of a K-12 AI curriculum be bracketed or removed from the next iteration of the Roadmap. At present, it is unclear what an AI-centric curriculum for grades K-5 would look like, as opposed to a curriculum that simply emphasizes STEAM. We also are unaware of any demand or support for this concept either within or outside the computing professions. If this issue remains in the text, however, we urge that it be reworked to make clear whether the Roadmap calls for educating K-12 students *with* artificial intelligence, *about* artificial intelligence, or potentially both.⁷

⁶ Many diverse approaches and conclusions are possible. For example, without here endorsing any view, we note that the Roadmap could state that AGI is a misleading and counterproductive idea and/or it could aver that ethical issues raised by AGI warrant deep exploration and present an excellent opportunity for computer scientists to engage in ethical discussions with others outside the field. (We note that specialists in fields like philosophy, psychology, neuroscience, law, and sociology will be critical to consult and engage in this context).

⁷ On the topic of using artificial intelligence to educate K-12 students, if such a goal is to be expressed in the Roadmap, then it will need to be significantly fleshed out and include discussions of educational theory, pedagogy, and even the economics of substituting computational educators for human ones. Building this into the Roadmap will require a significantly broader authorship and review committee, something which we also have suggested.

5. Expressly acknowledge that agency and accountability for the development and effects of AI systems reside in the humans behind AI, not the systems themselves, and articulate the need for many forms of robust research oversight, including both academic Institutional Review Boards and their industry equivalents.

The Roadmap as written may be read to make the troubling and misleading assumption that agency and accountability can reside in AI systems themselves. The Committee, and we expect many CCC colleagues, feel strongly that accountability for the determinations and effects of AI systems resides solely in human beings. Accordingly, we recommend that the Roadmap make this point clearly, repeatedly and forcefully.

Similarly, we believe the Roadmap could be improved by clarifying what forms of prospective, continuous, and retrospective oversight would best accelerate the development of AI governance norms and legal structures, thus helping to ensure the safety, security, reliability, and fairness (to mention a few) of AI-controlled or informed systems. Robust research is particularly needed on how to monitor and regulate critical systems that evolve quickly (relative to current regulatory processes) to cope with both inherently dynamic machine learning and agile systems development.⁸ In practice, potential problems will need to be rapidly identified and resolved to ensure that the system remains in a safe state.

We further recommend that the Roadmap expand its discussion of research oversight beyond academic Institutional Review Boards to directly address effective AI oversight by: government agencies;⁹ major accounting firms;¹⁰ insurance companies (which are evolving models to enable them to insure medical, manufacturing, and transportation systems that depend on AI); industry volunteer groups and trade associations, as well as civil society organizations.¹¹

For example, a number of vignettes in the Roadmap discuss research interactions with human subjects, but the accompanying research challenges are silent regarding how to manage user advocacy or stakeholder oversight beyond the initial research stage. The discussion of research interactions with human subjects only considers such interactions within the context of academia in the US.¹²

⁸ Retrospective analysis of failures may be the best (though certainly is not the only) means of maximizing system improvement. The Roadmap thus might productively discuss noted failures of AI systems that have resulted in problems for individuals, organizations, or society as a whole.

⁹ NTSB, FAA, FDA, FTC, FCC all are moving to develop and adopt some form of AI system oversight.

¹⁰ Ernst & Young, Deloitte, KPMG and other leading firms all are working to expand their corporate auditing service offerings and capabilities to include AI systems.

¹¹ The Partnership on AI, the Public Voice, and the OECD (among others) have promulgated guidelines for AI design, development, and implementation.

¹² In US academia, research involving humans is typically regulated by Institutional Review Boards. In the Roadmap, the only reference to IRBs occurs in the Accelerate Scientific Discovery vignette discussion (*Roadmap*, p. 39). Aside from being clearer on the particular relevance of IRBs in the context of this vignette, we suggest that much stronger ethics awareness information also be incorporated into other germane examples. For example, the Social Justice vignette describes a tenant eviction situation (*Roadmap*, p. 43), but only from the tenant's point of

Specifically, the Committee believes that the Roadmap can and should consider and articulate a plan for some type of ethics board-like oversight throughout the document. Currently, many Institutional Review Boards concern themselves with risks to the welfare of individuals or groups who are the subjects of research. Consideration of broader societal impacts during advocate/stakeholder reviews would help researchers and practitioners be more mindful of the effects of their work throughout their projects.¹³

6. Highlight broadly that research regarding the ethical challenges raised in and by the development and deployment of AI is particularly important, while also specifically noting that select issues merit intensive attention, such as: the use of AI in weapons systems and health care, algorithmic bias, and the disparate impact of AI's application on various classes of individuals.

We look forward to contributing to the additional analytical and drafting work required to address the ethical issues referenced above, and perhaps others, in the next iteration of the Roadmap. For the present, we believe that the Roadmap would benefit from a discussion of research opportunities for increasing trust, responsibility, and accuracy through AI system auditing.¹⁴

Further, because autonomous weapons systems have been a point of significant discussion within both the academy and industry worldwide, we also wish to highlight our conviction at this time that the Roadmap would benefit from specifically referencing the creation of autonomous weapons systems in the sections that discuss ethical issues raised by the development and deployment of AI.¹⁵

view. Presumably, there will be other parallel AI systems, e.g., ones that warn the landlord about potential legal infractions, or which trigger a legal aid worker's outreach before an eviction order is issued. These introduce other, obvious research challenges unaddressed by the Roadmap, e.g., should these systems be competing or collaborating? Who will oversee that process?

¹³ For example, the Roadmap's discussion of modeling mental and emotional states (*Roadmap*, p. 41) could benefit from enhanced commentary regarding the societal impact of AI and discussion how the Roadmap fits into the framework developed by the Menlo Roadmap.

¹⁴ For example, the IEEE Ethically-Aligned Design report (*Ethically Aligned Design*, IEEE Standards Organization, 2018, https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead_v2.pdf) and the EU Guidelines for AI (*Ethics guidelines for trustworthy AI*, European Commission Digital Single Market Report, 8 April 2019, <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>) specifically define and ground their work on the notion of "trustworthy AI". The Roadmap should expand its sole mention of auditing at page 112 ("AI system metrics would enable auditing and tracking compliance. . . ." into a full discussion of what constitutes trustworthy AI, why it is important, and how it can be achieved (or at least how research agendas regarding these issues may be formulated).

¹⁵ Similarly, in our view, the sections of the Roadmap discussing ethics and ethical issues could be productively enhanced by full discussion of research agendas for computational ethics. Presently, the Roadmap simply notes that there are ethical issues to be addressed.

7. *Strongly articulate the importance of integrating social and ethical agendas into computer science research and development from the outset of all such efforts (particularly those AI-related), rather than addressing such agendas only after product failure.*

We believe that the Roadmap would benefit by proposing an integrated social and ethical research agenda. Increasingly, academics, researchers and users of AI acknowledge that there are many deep and pervasive ethical issues raised by the technology, from the introduction of bias to the replacement of humans with computers.

The roadmap would be strengthened both by articulating the opportunities to include discussion of social and ethical issues, and by proposing research on these issues as well.

One set of ethical issue that the roadmap discusses, but only at times and tangentially, is matters of diversity and inclusion. We concur that it is and will remain important for the Roadmap to discuss the impacts of AI of all kinds on underrepresented and underserved populations. We believe, however, that this would best be accomplished by addressing AI's impact on many of the specific historically disadvantaged groups that will use and benefit from -- or be harmed by -- these systems including, for example, people who are illiterate, non-English speaking, disabled, and the infirm.

The Roadmap should bring such topics to the forefront, specifically by expanding its discussion regarding the need for research on the impact of AI on disadvantaged and marginalized populations. CCC should accomplish this, we believe, by expanding authorship and review of the Roadmap to include representatives of both the academic and non-academic research communities.

8. *Focus overtly on matters of diversity and inclusion, specifically by deepening discussions of the impact of AI on disadvantaged and marginalized populations, and expanding authorship and review of the Roadmap to include representatives of both the academic and non-academic research communities.*

Specifically, we suggest that the Roadmap affirmatively encourage researchers to seek ways to greatly increase diversity and inclusion in their own work. We also recommend that this goal be promoted with respect to production of the Roadmap itself by assuring that the text of the document ultimately is the product of a diverse and inclusive brace of experts,¹⁶ and by broadly soliciting public comment on it. We also urge that the next draft of the Roadmap be made available for public review over a longer comment period than provided for its first draft.¹⁷

¹⁶ Indeed, we note that the Roadmap's authoring committee is mainly comprised of representatives of U.S. research universities. Consequently, we are concerned that it may not fully apprehend or address the needs of AI researchers and users outside of those environs, including those in the industries referenced in the Roadmap. We thus recommend that the Roadmap's committee be expanded to include representatives from both industry and government research communities.

¹⁷ For example, the Roadmap text states: "Science requires a high standard of ethics and must adhere to IRB regulations and thus, the system must be able to address ethical concerns." (*Roadmap*, p. 40). Institutional Review Board (IRB) regulations only apply, however, to researchers at organizations that have such boards (primarily universities) or who receive research funds from the U.S. Government.

9. Include a new “AI and Security” section in the Roadmap.

Furthermore, although there is a discussion that AI is relevant to national security, there is little in the Roadmap about security and safety standards for AI applications. AI is seen as contributing to national security, but the Roadmap does not acknowledge or elaborate upon how reliance on AI could become a vulnerability.

Finally, while we are pleased to see mentions of adversarial issues in machine learning as something to be defended against, the Roadmap also does not mention whether developing attacks against statistical machine learning systems—offensive AI weapons—is included in this research agenda, or if the development of weapons is “off the table.” Clarification of this point would be useful.