The Computing Community Consortium's Response to RFI on National Priorities for Artificial Intelligence

July 7, 2023

Written by: *Justin Biddle (Georgia Institute of Technology), Elizabeth Bradley (University of Colorado Boulder), David Danks (University of California San Diego), Maria Gini (University of Minnesota), Ashok K Goel (Georgia Tech), Arjun Guha (Northeastern University), Madeline Hunter (Computing Community Consortium), David Jensen (University of Massachusetts Amherst), Sven Koenig (University of Southern California), Chandra Krintz (University of California-Santa Barbara), Daniel Lopresti (Lehigh University), Amy Mcgovern (University of Oklahoma), Daniel Molzahn (Georgia Institute of Technology), Rajmohan Rajaraman (Northeastern University), Aarti Singh (Carnegie Mellon University, (Matthew Turk (Toyota Technological Institute at Chicago),and Holly Yanco (University of Massachusetts Lowell).*

This response is from the Computing Research Association (CRA)'s Computing Community Consortium (CCC). CRA is an association of nearly 250 North American computing research organizations, both academic and industrial, and partners from the professional societies. The mission of the CCC is to bring together the computing research community to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges.

Introduction

AI technologies are progressing rapidly, with new tools and capabilities appearing daily, startup companies proliferating, and large tech companies pouring more and more resources into developing large models and infrastructure. Both the myriad potential benefits and the significant risks of these technologies have been discussed widely in recent years. In considering a strategy to harness potentially transformative AI benefits while mitigating the risks, a critical underpinning will be the ongoing engagement of the academic community in both research and education. These capabilities sprang from basic research and continued basic research will be required to explore new directions, seed innovation, and look beyond short-term applications. Higher education is also critically important to educate the growing cadre of AI scientists, engineers, and developers who will implement the coming generations of AI systems. Incentives must be created to ensure that government and industry continue to engage academia, including the sharing of large, expensive models and data needed to push the technologies forward.

This response includes answers to questions 1, 3, 5, 6, 8, 11, 12, 13, 14, 15, 16, 18, 22, 25 and 29. In our responses we try to stress the necessity of continued investments and emphasis on basic academic research and education. Academic research provides a unique space to explore critical, interdisciplinary questions unbiased by financial interests. Detailed responses are below.

1. What specific measures—such as standards, regulations, investments, and improved trust and safety practices—are needed to ensure that AI systems are designed, developed, and deployed in a manner that protects people's rights and safety? Which specific entities should develop and implement these measures?

Any broad AI policy should recognize the necessarily contextual nature of such policy. There are many different individual technologies that can be considered "AI technologies," and more such technologies will be developed in the future. The capabilities, benefits, and risks of AI technologies depend very much on the specific application, the specific technology, the way in which the technology is incorporated into a usable system, and the policies and procedures surrounding the use of that system with respect to that application. Recognizing that general "AI policies" will be few and far between, general big picture goals are:

1. All AI systems should be able to explain their information processing and justify their decisions and conclusions such that humans are assured of their reliability.
2. All machine learning systems must provide adequate information about, in some cases including access to, the data they are trained on.
3. All interactive systems must be transparent so that human users have a mental model of their functioning.

These goals are much easier said than done and require a complex and interactive web of government, industry and academic standards and regulations. Questions pertaining to measures that protect people's rights and safety usually focus on what can be accomplished through government regulations and industry standards. While such measures are undoubtedly important, the key role that academic research plays in the invention, refinement, evaluation, and promotion of new AI technologies receives less attention but is equally significant. Within the sphere of academic research, community standards are the driving consideration; they determine which ideas take hold and which wither on the vine. Academia also occupies a special position not only on the creation side, but also by providing a valuable counterbalance – an independent voice to critique what might be wrong, unfair, and dangerous with new technologies. The challenge is that AI research is spread across a number of distinct communities that function autonomously, and their standards have evolved independently: standards are known and enforced by members of each community, but not always formalized or recorded. For this reason, in parallel with efforts to develop government regulations and industry standards, we believe there is a need for a unifying effort across academic research to develop and deploy uniform community standards that protect the rights and safety of individuals where AI systems are used. It is also critically important to continue the funding of unfettered, "blue sky" research to develop new AI

technologies as well as work to view it with a critical eye. This is the basic framing we use for the rest of our responses to the questions raised in this RFI.

3. Are there forms of voluntary or mandatory oversight of AI systems that would help mitigate risk? Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?

Some mandatory oversight is essential. Europe has recently approved regulations for how to use AI, so AI companies who want to do business in Europe will have to obey the European rules. The proposed European AI Act uses a tiered structure, where unacceptable risk applications would be banned, high-risk would be regulated, and low-risk would just require disclosure. Other countries are working on their own rules, so it is possible that rules will differ by country.

An analogy that could be considered is the creation of an international convention, such as the Biological Weapons Convention, to regulate uses of AI and settle disputes. The UNESCO "Recommendations on the Ethics of AI," which have been adopted by all member states, provide a foundation for ethics in AI but do not address oversight, and they are meant for ethics, not other issues, such as mandatory requirements.

5. How can AI, including large language models, be used to generate and maintain more secure software and hardware, including software code incorporating best practices in design, coding and post deployment vulnerabilities?

AI and large language models (LLMs) have a huge potential to generate and maintain secure software as well as incorporate best practices in software engineering. For instance, recent research in programming languages has shown that an LLM can be trained on feedback from traditional analysis tools and then generalize to scenarios where traditional analysis fails, e.g., when the code is incomplete. Similarly, AI methods can be effective in suggesting fixes for post-deployment vulnerabilities. However, this has to be done carefully. Recent research has demonstrated that an improper training of code-generating LLMs can end up amplifying security vulnerabilities. Presently, there is very little research in this space, and the current products that do incorporate AI-powered debugging warn the user to separately check for software vulnerabilities. Research in academia and academia-industry collaborations are both critical to developing robust frameworks for training and inference of AI models for code, leading to secure software generation and maintenance.

6. How can AI rapidly identify cyber vulnerabilities in existing critical infrastructure systems and accelerate addressing them?

AI can detect cyberattack patterns, suspicious network activity, and malware much more effectively and quickly than people can, continuously updating its vast knowledge of cybersecurity risks and threats from within an organization, a region, and around the world. AI systems can analyze relationships among various factors – such as malicious files, suspicious IP

addresses, known adversarial organizations, and time-based events – to enable fast decision-making about complex events and potential responses and timely responses to security breaches. More specifically, AI systems can improve cyber security capabilities in many ways, including the following:

- Automated vulnerability scanning of critical infrastructure to identify potential weaknesses,
- Monitoring threat intelligence data from a wide range of sources and identify potential system vulnerabilities,
- Identify unusual patterns of activity that could indicate a potential cyber vulnerability or an ongoing attack and respond rapidly, and
- Identify areas of infrastructure that are more prone to vulnerabilities and propose remediation efforts.

In addition to monitoring, detection, and thwarting actions against a specific infrastructure system, AI systems can orchestrate information, security processes, and responses to cyber vulnerabilities across organizations, sectors, and beyond, including measures (such as anonymizing and aggregating data) to keep information private and safe. These capabilities can enable rapid, collective responses that help to combat attacks across multiple critical infrastructure systems.

8. How does AI affect the United States' commitment to cut greenhouse gasses by 50–52% by 2030, and the Administration's objective of net-zero greenhouse gas emissions no later than 2050? How does it affect other aspects of environmental quality?

Achieving the goal of reducing greenhouse gas (GHGs) emissions by this deadline will require that we target *all* of the leading sources of emissions including those from industry, transportation, agriculture, energy generation, and others [1]. Reducing GHGs within and across these sectors will require new multidisciplinary advances that leverage AI systems to manage the associated complexity. AI will play a key role in aiding and expediting the processing, aggregation, and analysis of a vast diversity and volume of data by automatically learning from this data and identifying opportunities for optimization that humans miss. By doing so, AI will make it possible to extract actionable insights from this data for use by humans in decision making as well as to facilitate data-driven automation, actuation, and control that will enable us to better monitor, predict, and adapt existing operations in ways that reduce emissions in cost effective ways.

AI can contribute to individual sectors in multiple ways, such as with:

- AI-enabled materials science (e.g., investigation and discovery of novel materials for energy storage),
- Planning, optimization, and decision support for production, distribution, and consumption of energy,

- Spatiotemporal planning strategies to optimize flows in transportation networks (traffic, goods, people, vehicles, pollution, energy, etc.),
- Algorithms that leverage agriculture sensor data, together with real-time information about economic factors and transportation networks, to inform planning and risk assessment, and
- Improving our understanding and prediction of high-impact weather events, such as flooding, wildfires, and severe storms and how these are changing as a function of climate change.

However, leveraging AI more aggressively to address these challenges will itself grow the carbon footprint associated with AI due to increased computation, communication, and storage. Therefore, our end-to-end solutions must account for the energy use and efficiency of the computing systems that bring them to fruition, and more research is needed in this direction.

[1] https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions

11. How can the United States work with international partners, including low- and middle-income countries, to ensure that AI advances democratic values and to ensure that potential harms from AI do not disproportionately fall on global populations that have been historically underserved?

AI is becoming so ubiquitous that everyone, everywhere has the potential to benefit, and conversely, everyone has the potential to be harmed. Because of the huge economic advantages they hold, entities that operate in and target audiences from wealthy economies lead in the development and deployment of AI systems. Those from low- and middle-income countries are excluded from receiving the benefits because they are outside the "core demographic." Worse, perhaps, they may be exploited as cheap labor for developing and improving AI tools not being built for their benefit, but for the benefit of others. It will be important to develop international partnerships aimed at leveling the playing field and treating all potential users of AI as equals, even when this is difficult to justify from a purely monetary standpoint. Beyond the regulations and standards that are needed, we again take note of the role academia can play. Most AI research communities are richly international, as are many institutions of higher learning. Even so, those who can participate still represent only the wealthiest of the world's population. Investments must be made to further diversity both education and research to reach those from low- and middle-income countries. We will also again note that community standards within AI research communities play an important role in the way new AI systems are conceived, developed, evaluated, and critiqued.

12. What additional considerations or measures are needed to assure that AI mitigates algorithmic discrimination, advances equal opportunity, and promotes positive outcomes for all, especially when developed and used in specific domains (*e.g.*, in health and human services, in hiring and employment practices, in transportation)?

A key consideration is the development of effective practice recommendations for engaging stakeholders who are potentially affected by AI systems and incorporating their needs and interests into the design of AI systems. This requires a participatory design that elicits, analyzes, and incorporates feedback from stakeholders regarding domain-specific sources of biases, design, and effectiveness of fairness metrics. It is important that the perspectives of impacted stakeholders – especially those from under-served communities – be included in the early stages of the development of AI systems, as early-stage design decisions can become entrenched and difficult to correct downstream.

The Computing Community Consortium recently held a workshop on Community-Driven Approaches to Research in Technology & Society. The focus of this visioning activity was to enable conversations between computing researchers and underrepresented communities that are impacted by artificial intelligence systems. This included active participation from a broad range of stakeholders from the disabled community, minority racial groups and from underrepresented classes of workers affected by AI systems, such as delivery drivers and bail navigators. The workshop also focused on outlining how community partners and researchers can effectively and ethically work together to conduct community-driven research. The community workshop report will be released in the next few months.

A specific example is the need for ethical consideration when AI and advanced optimization tools are applied for infrastructure systems like electric power grids. Taking advantage of new sensing, communication, and so-called "grid-edge" capabilities, researchers are deploying advanced algorithmic tools for designing and operating infrastructure systems which have the potential to significantly improve system-wide outcomes in terms of resilience, reliability, and cost. However, this effort, like many others, may not be embedding fairness considerations into their development. The NSF Artificial Intelligence (AI) Research Institute for Advances in Optimization (AI4OPT) is working on a project for power system reinforcement (adding batteries, solar panels, line undergrounding, etc.) so that power outages are distributed more fairly throughout the system when they occur [1]. Another example is an operational problem in the context of severe wildfire conditions [2].

Additionally, mitigating algorithmic bias requires developing practice recommendations for both data documentation – including documenting biases in datasets for training ML systems – and for mitigating biases in existing training data sets. Data needs to be vetted to make sure it is well understood and accompanied with a required description of circumstances/context involved in data collection should be documented to identify potential sources of bias. Each AI tool should be accompanied with a clear list of data (and what part of a specific dataset) that was used for training. Additionally, it could be helpful to develop benchmarks to ensure data is FAIR (findability, accessibility, interoperability, and reusability) to the extent possible (e.g., providing meta-data if data needs to be protected). There is a reason to be particularly wary of synthetically generated data. There is a big draw towards using synthetically generated data with advances in AI methods and synthetic data generation's heightened ability to respect privacy constraints;

however, these synthetic data sets come with all the same algorithmic biases and should be understood and documented.

[1] https://www.ai4opt.org/

[2] https://ieeexplore.ieee.org/abstract/document/9993295

13. How might existing laws and policies be updated to account for inequitable impacts from AI systems? For example, how might existing laws and policies be updated to account for the use of generative AI to create and disseminate non-consensual, sexualized content?

One general issue is that existing laws were often written with assumptions about the scale or breadth of potential impacts and harms. However, present-day and near-future AI systems are able to perform vastly more decisions and actions, including the generation of content. Moreover, these decisions can have far-reaching impacts that outstrip any expectations of the original drafters of the laws. We encourage the administration to carefully examine whether existing laws aimed to prevent harms are appropriate for the scale of harms that are possible from AI systems.

A second issue is the way in which AI systems can be used to target particular groups or communities without any explicit (and thereby actionable) decision to intent by a human. Current laws are not well-designed to control, reduce, or mitigate these types of algorithm-driven harms, as there is typically no individual, or organization that clearly engaged in illegal behaviors. Hence, we encourage a close examination of current laws and regulations to ensure that responsibility and liability for harms cannot be avoided merely by interposing an algorithm in the middle of the process.

14. How can AI be used to strengthen civic engagement and improve interactions between people and their government?

AI has the potential to greatly assist the public in their interactions with government and the services offered. While virtual assistants and chatbots are already being used at all levels of the US government, future improved refinements of large language models can be effective at answering questions, explaining rules and regulations, and directing people to suitable resources. The use of AI in government needs great care, however: automating the decision-making process carries huge risks owing to current AI tools' susceptibility to bias and lack of transparency and explainability. There is also a balance to be considered between AI tools provided by a government agency, versus a personal AI tool that is known better and trusted more by an individual citizen; freedom of choice should be supported. More novel and promising applications of AI in government include technologies for participatory governance, which will enable communities to participate in decision-making by experimenting with models and contributing to the design process in matters such as public transportation and urban planning.

Turning back to our original framing, we note that academic research can help answer questions like these. In this case, interdisciplinary partnerships between AI researchers, social scientists, and specialists in fields such as education will play a vital role. There is history to support the effectiveness of such collaborations, but the speed with which AI is developing and the ubiquity of its impact creates special challenges.

15. What are the key challenges posed to democracy by AI systems? How should the United States address the challenges that AI-generated content poses to the information ecosystem, education, electoral process, participatory policymaking, and other key aspects of democracy?

AI is a disruptive technology that poses many challenges for individuals, groups, and society as a whole. Some key challenges/concerns are privacy, bias and discrimination, lack of transparency, job displacement and economic inequality, manipulation and disinformation, and the concentration of power into a small number of entities. A multifaceted approach is needed to provide an environment in which society can benefit from the power and promise of AI technologies while minimizing the negative and destructive efforts of malicious players and unintended consequences.

As a more specific example, two critical challenges are: (1) improved micro-targeting of messages and content (whether factual, misinformation, or disinformation); and (2) rapid generation of synthetic content that humans cannot distinguish from truthful content. In both cases, these AI-based capabilities are likely to have harmful effects, both first-order (e.g., people believe falsehoods or operate in echo chambers) and second-order (e.g., people disengage from, or become highly cynical about, the political process). Significant research and development is needed to further develop AI technologies that can detect and mitigate the harmful effects of AI-generated content, including content provenance and verification, detection of AI-faked and manipulated media, and understanding the magnification effects of communication and social media applied to misinformation that has been generated by AI.

Unfortunately, there are significant limits on the ability of the U.S. government to directly address these issues, and so we suggest that much of the focus should be on education, both civic and technical, alongside efforts to increase civic participation, including (but not limited to) electoral and other political processes. The U.S. should invest significant effort in constructing effective regulatory frameworks that address these key challenges, including ethical guidelines and standards for the development and use of AI systems in various contexts. The country must also invest in R&D in AI technologies that can minimize, detect, and mitigate harmful effects of the technologies. Addressing AI challenges requires incentivizing collaboration among government, industry, and academia, promoting transparency, accountability, and sharing of systems, data, and vulnerabilities. General AI literacy initiatives in education will also be important, as will campaigns that raise public awareness of AI systems and their challenges.

We encourage the administration to consider more direct responses – for example, the use of truth-in-advertising regulations or digital watermarking to help identify AI-generated content – but these should be done with full awareness of their limits.

16. What steps can the United States take to ensure that all individuals are equipped to interact with AI systems in their professional, personal, and civic lives?

It is critical that people learn to understand AI systems and the ecosystems in which they exist in ways that allow them to comprehend, benefit from, and respond to the technologies and the various issues – both technical and social – already impacting our lives. Key steps include:

- Education: Educational initiatives are needed to raise awareness and provide foundational understanding of AI concepts and applications. This will require integrating AI-related topics into school curricula at all levels, promoting AI and data literacy and raising awareness of the technology's benefits, risks, and ethical and legal considerations – including privacy, consent, fairness, transparency, accountability, accessibility, and user rights and values. Vocational training must also be considered to provide an AI-literate workforce at all levels.
- Access to tools and resources: It will be critical to ensure equitable access to AI technologies, tools, and resources – including software platforms, data, advanced large models, user-friendly and explainable interfaces. User-centered design practices should be employed to make AI accessible to users with varying levels of technical expertise.
- Partnerships: Collaboration among government, industry, academia, and other aspects of civil society is crucial to facilitate the development of shared systems and data, education and training initiatives, and other key resources.

18. How can the United States harness AI to improve the productivity and capabilities of American workers, while mitigating harmful impacts on workers?

AI-powered "copilot" tools have great potential to enhance the productivity and effectiveness of many work functions, augmenting workers in a wide range of tasks and domains, such as:

- Writing – going beyond spell- and grammar-checking, AI tools can be used to improve clarity and readability, expand relevant source material, in summarizing material, in translating, etc.
- Creative content – AI copilot tools can augment the creative processes of artists, musicians, designers, and scientists by offering suggestions, generating rough draft content, refining creations, and more.
- Financial analysis – Financial AI tools can analyze massive amounts of data and provide deep insights and recommendations for portfolio management, risk assessment, and other aspects of financial decision-making.

- Information – AI can extract and summarize information from large text documents, videos, and other sources, saving time and increasing the likelihood of finding needles in the large haystacks of data.
- Healthcare – AI copilot tools can assist healthcare professionals in diagnosing and treating patents, analyzing medical records and lab results, searching medical research for relevant cases, and so on – in general, supporting more accurate, efficient, and effective decisions and improved medical outcomes.
- Education – AI tools can provide personalized learning experiences that adapt to an individual student's experience, strengths, and limitations. It can also offer feedback on performance and support both students and teachers in providing more effective learning environments.
- Software development – AI-powered programming tools can allow developers to focus on the important task-level aspects of a problem and provide analysis and testing recommendations that reduce the likelihood of incorrect results, security problems, and software bugs.
- Legal work – AI tools will be increasingly useful for legal tasks such as analyzing large volumes of documents (contract analysis, case research, searching for applicable precedents), generating novel legal theories to explore, testing and refining legal arguments, and assessing legal compliance.

While the above examples identify a space for human-robot cooperation, the World Economic Forum states that an estimated 44% of jobs will be disrupted within the next five years [1]. These workers will need reskilling or upskilling. Many have families and jobs and thus they cannot move to places of education; online education offers a medium to take education to where they live and work. However, the quality of online education is uneven. To mitigate these negative impacts to workers we must use AI to enhance the proficiency of online education so that workers can learn the skills they need in a fast-changing workforce.

[1] https://www3.weforum.org/docs/WEF_Future_of_Jobs_2023.pdf

22. What new job opportunities will AI create? What measures should be taken to strengthen the AI workforce, to ensure that Americans from all backgrounds and regions have opportunities to pursue careers in AI, and otherwise to prepare American workers for jobs augmented or affected by AI?

The recent advances in deep learning, large language models, robotics, computer vision, and machine learning are transforming diverse industry domains including medicine, transportation, manufacturing, finance, agriculture, and defense. There is considerable uncertainty about the evolution of AI technologies and the nature and extent of their adoption in disciplines. Ongoing research on the future of work points to a transition period with possible job displacement across areas. Nevertheless, the potential for AI to spark new discoveries in most human endeavors is clear, and this will create opportunities for workers with broad knowledge of both the benefits

and downsides of AI, and for knowledge-workers proficient in AI tools. Therefore, an essential measure to strengthen the American workforce is AI education that ensures that Americans from all backgrounds and regions are knowledgeable about the applications, strengths, as well as limitations of AI.

The growth in opportunities for AI systems will create demand for professionals who can build, assess, train, calibrate, and optimize these systems for specific tasks and domains, and for many jobs related to their safe and effective deployment. These include required expertise in data collection and annotation, AI explainability, AI auditing, data and system security, user experience for AI systems, advanced robotics, AI business strategy, AI policy matters, and AI ethics.

Currently, most of the information that the public gains about AI is through news sources and social media, which often present a distorted view of the technology that emphasizes the extremes over a balanced perspective. Teaching AI and its numerous applications in elementary and secondary schools is critical to introduce students to an unbiased treatment of AI and prepare Americans for the jobs of tomorrow. Also important is investment in AI research and higher education to spur the development of new techniques and to train early-career engineers and scientists to work at the forefront of AI development.

25. How can Federal agencies use shared pools of resources, expertise, and lessons learned to better leverage AI in government?

Shared resources among government, industry, and academia are critical, as the perspectives and contributions of each group are different and vital. Shared data, models, platforms, and computation must be incentivized so that all groups have access to run, experiment with, and modify powerful, state-of-the-art AI systems. Large-scale AI research funding, such as the National Artificial Intelligence Research Institutes [1], can serve as hubs of expertise, training, best practices, and shared knowledge, data, and models.

Federal agencies should coordinate to provide information, advice, data, and funding to create and run such research and collaboration activities and to provide insight and direction to individual agencies and cross-agency activity. Agencies may also benefit from contributing R&D staff to such research projects who can be embedded in the research and development for a time and bring cutting-edge knowledge and practice back to the agencies.

[1] https://www.nsf.gov/cise/ai.jsp

29. Do you have any other comments that you would like to provide to inform the National AI Strategy that are not covered by the questions above?

We are very pleased to provide our responses to these questions, which we feel are right on target for many of the most important issues we face as a society. The coming large-scale deployment of newly-powerful AI systems will be incredibly disruptive: there is no avoiding

that. The main question is finding the right path to minimize the negative impacts and encourage approaches that benefit humanity and protect everyone's rights and safety. The challenge is so large that it cannot be solved purely through government regulations or industry standards, as important as they are. While academic research, much of it federally funded, has already played a major role in advancing AI, it is clear that many questions remain in transitioning from early work that interested a few specialists to the fielding of AI systems that impact everyone in society. Our institutions of higher learning are uniquely positioned to examine the hard research questions with a critical eye, and to provide the education necessary for us to continue to lead in developing and deploying the kinds of AI systems that meet our goals and aspirations as a society.