WHITE HOUSE OFFICE OF SCIENCE AND TECHNOLOGY POLICY

Request for Information: Preparing for the Future of Artificial Intelligence

Comments of the Computer & Communications Industry Association

The Computer & Communications Industry Association commends the White House's Office of Science and Technology Policy for its framing of this Request for Information on the future of Artificial Intelligence (AI). As the RFI notes, AI technologies offer great promise for new and innovative products and services, economic growth, and applications across society.

In discussing the potential benefits of artificial intelligence, it is helpful to sketch out the contours of the field. Today, AI refers to the technical discipline of making machines intelligent—computational systems that can respond to complex factors in a particular context to achieve some goal.

Rather than a theoretical technology relegated to science fiction, AI is a tool currently used by academics, engineers, and scientists worldwide. Present research and development is focused on the practical application of AI to existing problems, rather than the development of an artificial general intelligence commonly portrayed in science fiction. Machine learning is a related discipline that has direct relevance to AI's use as a problem-solving tool, because it enables systems to make inferences from large samples of data.

While AI has been developing rapidly in recent years, its continued progress and impact cannot be taken for granted. AI has the potential to transform healthcare, transportation, security, education, and more—but only if stakeholders work collectively to encourage its innovative potential.

1. The legal and governance implications of AI

Like any new technology, artificial intelligence and its practical applications can raise regulatory and legal questions. As AI technologies develop, so too will society's ability to manage their use and to determine areas of possible concern. The ultimate goal should be to avoid actual harms and promote innovation in and the use of AI.

The range of potential uses for artificial intelligence is enormous and cuts across sectors. No single regulatory solution will appropriately respond to all possible concerns. But AI-enabled platforms are not emerging in a regulatory vacuum. The data and activities associated with the most sensitive applications of artificial intelligence are already subject to the protections of existing rules, which cover areas including privacy, data security, energy, finance, and transportation. Each of these sectors has an expert agency with knowledge and tools available to ensure that any harms AI might pose are appropriately addressed. The government should convene these agencies and stakeholders before considering new regulation to properly apply the protections of existing rules.

If specific new rules are deemed necessary to respond to concerns about AI, policymakers should look toward principles-based guidelines where possible. Best practices developed through stakeholder consensus can help drive innovation while providing protection where necessary. As appropriate, the government should convene stakeholders to aide the development of industry-wide best practices and self-regulatory regimes for the various applications of AI.

2. The use of AI for public good

Artificial intelligence and machine learning can be used for the public good in a variety of fields, including healthcare, cybersecurity, and education.

The application of AI to healthcare problems will allow physicians to be more accurate, see more patients, and save more lives. AI can reduce human error by helping scientists and clinicians detect patterns in medical data, diagnose illnesses, and recommend treatments. Several startups around the country already use machine learning techniques and predictive analysis to provide personalized healthcare guidance to patients, improved follow-up care, and better identification of new pharmaceutical therapies. Using AI in healthcare improves the quality of care, lowers costs, and delivers better outcomes.

Cybersecurity, another data-driven field, is similarly primed for AI-enabled growth. Intelligent algorithms are beginning to form the core of real-time threat prediction, detection, and response on secure networks in the event of a cyberattack. For example, machine learning enables systems to understand normal user and network behaviors, for later use in identifying deviations that signal possible intrusions. Similar tools deployed by information sharing and analysis organizations can coordinate to detect fraud, combat breaches, and reduce identity theft across sectors and regions.

Significant social benefits will result from the application of AI to education. Smarter software can help teachers customize lesson plans based on individual students' needs and automate basic activities. Students will benefit from educational software that adapts to different learning styles and paces of study, which can also facilitate remote instruction.

4. The social and economic implications of AI

    a. Economic implications

Artificial intelligence can lead to efficiencies and productivity improvements across the economy. AI-enabled modeling software can help analyze data, manage records, automate information acquisition, optimize logistics, and produce valuable insights about markets. The Analysis Group recently estimated that AI could have an aggregate economic impact of $1.49 trillion to $2.95 trillion over the next ten years.

The cumulative economic effects of advances in artificial intelligence and deep learning are likely to be positive, both in terms of labor participation and labor productivity, as proven by many prior technological innovations. Although the concern of "AI replacing humans" has received significant attention, AI does not mean automation. A more accurate representation of the effects of AI, particularly in the short and medium term, is a future in which deep learning augments human labor to increase workforce productivity and help create new jobs.

These productivity boons will be particularly important for small businesses. Smart platforms can boost economic activity by large numbers of small enterprises by allowing them to intelligently scale their businesses and empower their employees through smarter tools.

    b. Social implications: avoiding discrimination

AI systems that help make decisions based on complex factors and data sets can raise concerns about unfair or discriminatory outcomes. These outcomes might result from design

choices or biases inherent in the data used to condition an intelligent system. If potential sources of bias are not unaccounted for, actual harms can result.

But well-designed AI systems can also help avoid discrimination in areas where it is unintentionally present. For example, present professional hiring practices can sometimes lead to unconsciously biased results. A number of new startups are helping to incorporate machine learning and automation into hiring processes. By using employers' own data and publicly available information to suggest candidates who might otherwise have been dismissed for reasons unrelated to qualification and fit, these startups are helping recruiters build more diverse and productive workplaces.

In seeking to avoid discrimination, policymakers should recognize that AI-enabled systems are simply tools. Existing laws that apply to sensitive areas like housing, finance, and employment already provide technology-neutral remedies for disparate impacts. It would be counterproductive to mandate human involvement in every AI system, since people often hold inherent biases. Regulators should instead aim to provide companies and consumers with tools to diagnose and prevent failures that might lead to discrimination.

Biased outcomes are also often the result a lack of quality data, which can negatively affect an otherwise well-intentioned machine learning protocol. To help rectify this, governments should facilitate the release of robust datasets that enable responsible analysis and use, especially in areas where AI systems are publicly deployed.

8. The specific steps that could be taken by the federal government, research institutes, universities, and philanthropies to encourage multi-disciplinary AI research

The government should enable policies that encourage research and development, foster the AI workforce, and promote public AI deployment.

    a. Encourage diversity in all aspects of AI development

It is imperative for innovation and AI development that communities be diverse and represent a broad set of backgrounds and experiences. Key companies advancing AI, such as Nvidia, Google, and Enlitic, were founded by immigrants. Immigrant academics have become some of the leading voices and advocates for AI in the U.S. and help shape its future workforce to take advantage of their expertise. The U.S. should increase the availability of H-1B visas to further capitalize on this worldwide talent pool.

    b. Invest more resources in STEM education

The government, universities and research institutes should prioritize the value proposition and flexibility of STEM disciplines when recruiting individuals, as these skills directly translate to improved AI research and development. Examples include the White House's recently launched Computer Science for All initiative, which enables students to develop computational thinking skills early, and the Department of Labor's the TechHire program, which provides federal funding for accelerated talent pipelines in STEM-focused sectors.

    c. Support internal government expertise in technology

The government should continue to expand its technical capabilities through programs like the U.S. Digital Service and 18F. Every agency will be better positioned to leverage AI technologies

for complex problems in their respective domains if they house experts in computer science and technology.

Similarly, the Presidential Innovation Fellows program aims to connect innovative thinkers with relevant government agencies and civil servants. Fellows bring expert knowledge and practices into the government to address some of the nation's biggest challenges at the convergence of technology, policy, and process. This collaborative, user-centric approach will be essential for implementing AI-based solutions across the federal government.

    d.   Leverage global innovation networks

The U.S. should support pro-innovation legal regimes abroad. It should continue to partner with other innovative countries to share resources and advance areas for cooperative growth. Concurrently, these countries should also establish partnerships with developing nations, which have proven that innovation-driven growth is no longer the prerogative of high-income countries and have increasingly designed policies to increase their innovation capacity.

Progress in artificial intelligence is the product of international collaboration. Copyright is one field in which the U.S. can promote pro-innovation frameworks. Machine learning in particular is dependent on balanced copyright laws that promote innovation. Machine learning generally requires the analysis of large samples of data and information to condition an intelligent algorithm, the availability of which may be restricted by copyright regulations in certain countries. In the United States, established limitations and exceptions to copyright, such as fair use, enable access to non-expressive use of works for innovative purposes. However, U.S. companies, especially startups and small businesses, may face anticompetitive restrictions in other countries, which the U.S. should work to address with its partners.

    e.   Share data and support AI research

Private companies often underinvest in research and development since return on investment for experimentation is uncertain. In the past, federally-funded research has been the catalyst for many of today's AI technologies. Today, the government provides 60% of the funding for basic research in AI. NSF has spent $200 million thus far on AI, and DoD supports research with the Machine Reading and Mind's Eye projects as part of its annual $250 million budget on big data. As AI matures, the government must continue to budget for basic research on machine learning and emerging AI technologies.

A key part of the evolution of AI technology has been the development of technology-powered platforms and services that bring value to the mass market. Prioritizing research and AI projects within and in collaboration with government agencies can foster greater academic participation and industry growth. Many advances in machine learning have been products of research projects largely funded by DARPA, which offers cash prizes to innovators who successfully complete challenges in fields like robotics. Expert agencies could foster collaborations between machine learning experts and domain experts in fields like medicine, science, and business. This approach to research increases transparency and creates informed strategies that benefit future developments.

In addition, some research communities can have insular technical conferences. Community leaders, the federal government, and philanthropies could all provide support for collaborative venues through which machine learning and domain experts can interact.

Finally, the federal government should facilitate the responsible analysis and use of AI systems through the release of accurate and robust datasets. The ImageNet visual database and image classification challenge have helped spur the recent commercial deployment of deep learning algorithms. The Department of Commerce has been particularly active in advocating for open data initiatives and making datasets available to businesses. Encouraging the creation and curation of new or better datasets in a variety of application areas will help further drive machine learning to society's benefit.