Microsoft

The Role of AI in National Security

Dr. William Chappell, Chief Technology Officer and Vice President
Strategic Missions and Technology
Microsoft

Introduction

Artificial intelligence (AI) is transforming every aspect of human life, from health care to education, from entertainment to commerce. AI has the potential to enhance the security and prosperity of the United States and its allies, as well as to promote the values of democracy and human rights around the world. However, AI also poses significant challenges and risks, such as ethical dilemmas, social impacts, and geopolitical competition. Therefore, it is essential that the development and deployment of AI is guided by responsible principles and practices that ensure the safety, reliability, and accountability of AI systems.

Microsoft is committed to advancing responsible AI across the industry and society. We have adopted six principles to guide our AI work—fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability—that we put into practice through implementation of our Responsible AI Standard. We also support the development of common standards and norms for AI that reflect these principles and protect the rights and dignity of all people. We believe that these principles are aligned with the Department of Defense's (DoD) responsible AI policies, which aim to ensure that the U.S. military uses AI in a lawful, ethical, and effective manner.

Every country capable of developing and deploying generative AI is doing so today. Some are doing it through massive state-level investments. In the United States, our competitive advantage lies in our combination of private sector innovation, including on AI technologies and responsible AI techniques; increasing government focus on investments in safe and rights-respecting AI deployment; and the development and implementation of effective, globally interoperable policies for AI safety, security, and trust. Continuing to harness commercial innovation and government investment will be critical to maintaining U.S. leadership on AI.

While the U.S. government has some unique needs, many of the jobs that AI is best at – supporting decision-making, boosting productivity, and improving cybersecurity, to name a few – are largely the same as in the private sector. Our partner companies are embracing this new technology and seeing impressive results. The U.S. national security enterprise, on the other hand, needs a concerted effort, with resources and a sense of urgency, to let mission users rapidly experiment with the technology at scale to show how it can be effectively and responsibly harnessed. Responsible, mission-driven AI comes from active experimentation and evaluation, not avoiding usage.

Microsoft recommends that the Department of Defense and the intelligence communities establish dedicated programs with centralized funding to support a competitive process to support promising artificial intelligence uses for mission. These programs should bring in private sector AI developers to leverage commercial advances in models and safeguards, and by centralizing funding, the programs will improve the speed and efficiency of interacting with industry and accessing critical compute. These programs should be in addition to the several ongoing agency- and service-specific AI initiatives already underway.

The Importance of Maintaining the U.S. Edge in AI

The U.S. military is a global leader not only because of its superior technology and equipment, but also because of its highly trained and dedicated personnel, its rigorous standards and discipline, and its adherence to the principles of honor and integrity. The U.S. military has a noble mission of protecting the nation and its allies, and of promoting peace and stability around the world. The U.S. military is also a leader and a partner in the global community, working with other countries and organizations to address common challenges and threats, and to support humanitarian and development efforts.

As a society, our use of generative AI has greatly accelerated during the past year, and academia is beginning to release studies confirming what we've been thinking for some time: this technology dramatically helps people do their jobs better. A recent Harvard Business School/Boston Consulting Group study,[1] for instance, found that consultants using GPT-4 competed 12.2% more tasks, 25.5% faster, and produced results of 40% higher quality. Interestingly, this study also showed that generative AI can close the gap between lower and higher performers by providing higher gains for lower performers.

Just as in the commercial sector, AI holds the potential to make military personnel even better at their jobs. Other countries know this, and through generous funding and support for their industrial champions, are catching up with and challenging U.S. leadership on AI. We've been fortunate so far in that the U.S. has been a global leader in AI research and development, thanks to its strong scientific and technological base, its vibrant entrepreneurial ecosystem, and its democratic values. However, the U.S. cannot afford to rest on its laurels.

It is imperative that the U.S. maintains its edge on AI and remains a standard-setter for responsible and trustworthy development and deployment. This requires a sustained and coordinated effort by the government, the private sector, academia, and civil society, to continue to invest in AI research and innovation, to foster a diverse and skilled AI workforce, to promote a robust and ethical AI governance framework, and to collaborate with like-minded allies and partners on AI issues. By doing so, the U.S. can leverage AI to enhance its national security and competitiveness, as well as to advance its values and interests around the world.

In addition to the specific national security recommendations elsewhere in this paper, Microsoft is strongly supportive of the U.S. government efforts to establish and fully fund a National AI Research Resource (NAIRR). The NAIRR would provide opportunities for our public sector scientists, academics, and others who wouldn't otherwise have access to high-end compute to pursue

[1] Dell'Acqua, et al. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Havard Business School Working Paper, 24-013.

cutting-edge AI research. This continuous improvement of artificial intelligence will also support ever-improving national security tools, provided these advances are integrated into mission.

Responsible Development of AI

Microsoft has developed a comprehensive framework for Responsible AI, based on six principles: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. These principles guide Microsoft's AI development, deployment, and use across its products, services, and operations. Moreover, Microsoft has developed a Responsible AI Standard, which defines goals and requirements for implementing our six principles and is applicable across engineering and product groups. Our Office of Responsible AI, communities of Responsible AI Champions, and Responsible AI Council, led by our Vice-Chair and President Brad Smith and our Chief Technology Officer Kevin Scott, foster both bottom-up and top-down governance to support implementation of our Standard and strengthen accountability.

Microsoft also established a Sensitive Use Case Review Process, which is a cross-functional and multidisciplinary process that evaluates the potential risks and benefits of AI use cases that may have significant social, ethical, or human rights implications. The process involves experts from engineering, legal, policy, ethics, and human rights teams, who assess the use cases against the Responsible AI principles and Standard, as well as the applicable laws, regulations, and standards. The process also considers the potential impact on the stakeholders, such as customers, users, employees, and society at large. The process aims to ensure the robustness of processes to identify and mitigate the risks, and to provide guidance and recommendations on how to proceed with the use cases, or whether to decline them altogether.

The Department of Defense has also recognized the importance of responsible and ethical development of AI, and has issued several policy statements and documents to guide its AI activities, and we're proud to have been able to share our Responsible AI approach with the Department over the last several years. In particular, DoD's AI Strategy outlines the vision, goals, and objectives for leveraging AI to advance the national security and defense missions. The strategy also emphasizes the need to uphold the highest standards of ethical behavior and moral values in the use of AI. In 2019, the DoD adopted its AI Ethical Principles – responsible, equitable, traceable, reliable, and governable – that provide a framework for the DoD to design, develop, deploy, and use AI in a lawful and ethical manner, while ensuring the trust and confidence of the warfighters, partners, and the public. The DoD has also established the Responsible AI Committee, which oversees the implementation and oversight of the AI Ethical Principles across the DoD.

While AI development and deployment must be done responsibly, it can also be done with a sense of urgency. DoD has the tools and structures in place to drive a culture of innovation and experimentation that will get AI into the hands of mission users quickly and responsibly.

The Need for Experimentation and Innovation

To maintain its competitiveness and ready itself to take advantage of new opportunities, the United States must accelerate the adoption and integration of AI into its defense operations. AI can provide the U.S. military with significant advantages in terms of speed, accuracy, efficiency, and effectiveness, as well as new capabilities and opportunities for innovation and transformation. AI can also help the U.S. military address the emerging and complex challenges and threats in the 21st century, such as cyberattacks, disinformation, and hybrid warfare.

However, the adoption and integration of AI into the U.S. military is not without challenges and risks. AI systems are often complex, dynamic, and opaque, which can raise technical, ethical, and operational issues. For instance, how can the U.S. military ensure the reliability and safety of AI systems, especially in high-stakes and high-risk scenarios? How can the U.S. military ensure the compliance and alignment of AI systems with the laws of war and the rules of engagement? How can the U.S. military ensure the human oversight and control of AI systems, and the accountability and responsibility for their actions and outcomes?

These issues require careful and rigorous testing and evaluation of AI systems, as well as continuous monitoring and improvement of their performance and behavior. They also require a culture of experimentation and innovation within the DoD, which encourages and supports the exploration and discovery of new and novel ways of using AI to achieve the DoD's missions and objectives responsibly.

DoD needs to be actively working with AI to determine what applications will be most helpful for its mission, and to establish best practices to keep the technology human-centric. For instance, AI capabilities are critical to sorting through massive quantities of changing and mixed-form data, all of which exist in quantity at the Department. This ranges from back-office business systems to tactical networks and communications to defense platforms and weapons systems. Industry can supply these capabilities – and the Department should lend its technical and market expertise to explore use cases. Industry can also be a partner in developing the policies and concepts of operations that are both effective for the mission and consistent with our values.

Our experience over the last several years has shown us that users have to actively engage with AI to really understand how to use it. There's no shortcut for experience. Using it is the only way the most interesting use cases will be developed – and it's the best way to incorporate it into our everyday work. Critically, it's also the best way to identify safety, security, and trust issues and address them effectively.

So far, we have seen some government engagement with AI, but it pales in comparison to what we're seeing with our commercial customers, such as KPMG & PWC in finance, Epic in healthcare, and AT&T in telecommunications. These users have embraced the philosophy of growing – responsibly – with the technology, not waiting for some finished product to be delivered in the future. They are acting with urgency, because they know that waiting would only give their competitors an advantage. U.S. national security agencies are facing the same decision today, but we're not seeing this sense of urgency.

The U.S. national security enterprise needs two things–high level leadership, and a program to drive wide-spread experimentation by users–to capitalize on AI for mission. Where agency leadership exists, we're finding pockets of innovators eager to find ways to improve by using this technology. For instance, we have worked with the Defense Technical Information Center to personalize our models so that we have responsibly and carefully baked government information into the decision processes that they use. Unfortunately, these specific initiatives are too few and dispersed; more needs to be done, and with urgency. In general, potential innovators seem stifled as they look for “top-cover” from their leaders before they responsibly explore this exciting new technology.

In particular, we endorse the idea of creating a more aggressive "Replicator"-type program that would pursue six to twelve AI-enabled projects every year of a similar scale and ambition as Replicator, but across the scope of the Department’s responsibilities. This would provide an environment to allow mission users to experiment with AI across the Department and learn from those experiences. It could start in areas that are safe and not controversial. Then it could apply the experience gained to harder challenges. Overall, this program’s goal would be to rapidly develop and deploy AI solutions for the DoD's most pressing and urgent problems in a responsible and deliberate manner.

These programs can and should get underway immediately, even as the National Security Memorandum (NSM) required by the October 30 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence is being developed. Properly scoped experimentation, with appropriate guardrails, can give mission users a critical window to evaluate use cases and controls, and even to identify new issues that should be addressed by the NSM.

Finally, widespread AI adoption within the national security agencies will require a workforce with the required skills to implement this new technology. These skills will be largely the same that individuals across society are developing through new courses and educational opportunities. For instance, LinkedIn and Microsoft are offering the first professional certificate in Generative AI in the online learning marketplace to ensure widespread access to the skills and knowledge necessary to seamlessly incorporate this technology into individuals’ personal and professional lives. Since June, over 220,000 people in the U.S. have engaged with the content in the professional certificate and over 7,000 U.S.-based learners have completed the entire pathway and passed the exam, which has been accessible at no charge and will continue to be so through 2025 to ensure that this knowledge is readily available and accessible to all. Microsoft also recently hosted a series of virtual events to highlight the fundamentals of GenAI and explore how it can help people create, solve complex problems, and benefit their communities.

Conclusion

AI is a powerful tool for national security, and the U.S. must ensure that it stays ahead and sets the standards for ethical and reliable AI. We are at the start of this era of discovery and progress. AI is, today, the least intelligent that it will ever be. Look how far we’ve come over just the last year, and just try to imagine where we’ll be in a year.

To use this amazing new technology properly and responsibly, our national security community needs a stronger culture of experimentation and innovation, led from the top and properly funded. This will enable and support the quick and effective adoption and integration of AI into U.S. military operations. Microsoft is dedicated to working with the DoD and the wider national security community on advancing responsible AI for the good and the security of the nation and the world.