



Forum for American Leadership

Generative AI: National Security Opportunities and Risks

August 18, 2023

This paper is a product of the Forum for American Leadership's [Technology and National Security Innovation](#) Working Group. The primary authors of the paper were Jamil Jaffer and Luke Murry.

The Bottom Line

AI is moving at 21st-century speed. Generative AI is not a new field, but in early 2023 it exploded in public consciousness due to advancements in real-world applications and user accessibility. It took ChatGPT [just five days to reach 1 million active users](#). By comparison, Facebook took 10 months and Netflix 3.5 years to reach the same milestone. As policymakers grapple with how to handle this new environment, fostering continued innovation and being clear-eyed about the national security implications of generative AI will be critical.

What We're Talking About: Generative AI vs Large Language Models (LLMs)

- **Generative AI is a tool (e.g., algorithms) that can produce new content (text, video, audio, etc.) based on exposure to other sources or data.** It creates new content by applying math to existing data.
- **An LLM, such as ChatGPT, is a subset of generative AI models that focuses on creating new content in the form of text.** Other models, such as Stable Diffusion or OpenAI's DALL-E, can generate images and additional forms of content.

What to Keep an Eye Out For

- **Be ready for new players; expect big players.** Just as the internet browsers Prodigy, AOL, and Netscape Navigator gave way to Chrome, Safari, and Edge, do not assume that ChatGPT will outlast its competitors. Big players have an inherent advantage when it comes to distribution. For example, Google's Bard—its version of an LLM application—was released to billions of Google's existing users just months after OpenAI's GPT-4. Likewise, Meta has made the decision to open-source its Llama LLM, allowing a broad swath of developers to work not only on building applications but on the model itself.
- **The People's Republic of China (PRC) will seek to take the lead but will face challenges in producing a better product.** While the PRC has put massive amounts of human talent (mostly educated in U.S. higher education institutions) and government funding towards innovation, and though Chinese firms have the ability to leverage stolen intellectual property (IP), the Chinese Communist Party's (CCP) prioritization on censorship and control—and its fear of its own people—may actually inhibit its indigenous growth of highly capable generative AI. Indeed, Beijing [has already issued stringent requirements on LLMs](#) in its attempt to control information flow, and it would not be surprising to see more to come.

Recommendations

- 1. Focus on fostering innovation, not stifling it.** Generative AI is still a nascent industry, but one that is evolving at a pace [drastically faster](#) than other recent tech industry breakthroughs, such as the Internet or the iPhone. The technology is bound to change in ways that policymakers almost certainly cannot predict. The U.S. leads the world in software innovation writ large at least in part because anyone—from the high school dropout to the Fortune 100 company—can write code and, if it is good enough, commercialize it. Regulatory proposals that would require government approval for all new AI algorithms are full of pitfalls and likely to dramatically hamper innovation. The U.S. government should be far more worried about the government’s ability to stifle new technology than properly regulate it.
- 2. Use existing tools to address concerns.** Instead of developing a significant regulatory footprint, as European allies appear to be on the verge of doing, policymakers should first look to existing law to prosecute the malicious application of generative AI. Many potential negative applications, from hacking to fraud, are already illegal. To the extent there are gaps in existing law, policymakers should focus on developing targeted laws at specific harms. If security measures ought to be taken, the government should develop generalized frameworks and guidelines based on industry best practices, like the [NIST’s recent AI Risk Management Framework](#). On the international level, values matter. Policymakers should be on guard against an effort by Beijing to set global rules and intentionally engage with its allies and like-minded partners to, as much as possible, synchronize rules and regulatory frameworks in ways that undermine free expression and economic freedom. Likewise, the United States and its allies should come together and ensure that any legal or regulatory regimes to be put in place respect and promote the values we share in common and which have generated tremendous amounts of innovation historically.
- 3. Do not wait when it comes to promoting innovation.** The U.S. is [ahead of its adversaries](#) in developing generative AI, but they are not sitting around. Russia and China have a history of significant engineering accomplishments and the ability, particularly in the case of China as noted above, to pour tremendous quantities of human capability and government money towards industry programs. The earlier the U.S. adopts these innovations and promotes their development and use in both public and private sectors, the more time it will have to take advantage of its lead. And it will be easier for the U.S. to maintain its lead if American investors and companies don’t sell software and parts to—or invest money in—Chinese companies that are focused on AI-applications that would have a harmful impact on U.S. national security.
- 4. Application is key.** Geopolitical advantage does not always come to those who invent new technology, but rather to those who figure out how to apply it best. As when the Internet first broke onto the scene, it is going to take time to understand what generative AI is most useful for, from both a commercial and national security perspective. Rather than fighting the generative AI trend, the federal government should embrace it and seek to understand how to use it to their advantage. Over the short term that means immediately studying how to apply this new technology to their distinctive missions and working to implement

changes as rapidly as possible. Over the long term, every agency should hire the technical talent needed to ensure more effective implementation and use of generative AI.

5. **People matter.** If the U.S. government is going to take advantage of generative AI, it has to understand it. But [less than 1%](#) of new AI PhDs today land in government jobs. That has to change. From significant tuition assistance, to making it [easier for digital talent to perform meaningful work](#) related to their expertise, to offering at least marginally competitive compensation, the government needs to think in big, creative ways to do a better job attracting tech talent.
6. **The national security community, in particular, needs to put aside its risk aversion and better utilize AI tools.** Both the U.S. Department of Defense and the intelligence community need to experiment, tweak and incorporate generative AI in how they operate immediately. This does not mean allowing AI to act independently without a human in the decision loop. Rather, generative AI can be used to enhance, not replace, human decision-making. For example, a ‘ChatGPT for Analysts’ that is specifically designed for classified contexts and pre-trained with all the intelligence reports that have been written, as well as all the relevant open-source data, would help human analysts assess the probabilities of future events. From synthesizing intelligence to bringing timely data to the warfighter (e.g., “where is my adversary likely to go next?”), generative AI can make a real, substantive difference. The posture of the national security community should be to lean in, not shy away.
7. **Be prepared to respond if the PRC gets out of its own way and seeks to endorse and use government-approved LLM.** Beijing could very well incorporate a PRC government-approved LLM into its information warfare tactics, using it to suppress anything that makes the CCP look bad, promote anything that casts the U.S. in a negative light, and to control the global narrative by pushing it out through its network of state-owned and state-influenced enterprises. To that end, if the Chinese do create their own government-approved LLM, the U.S. government should be sure to prevent such an LLM from gaining a TikTok-like foothold domestically and should encourage allies to do the same.

Conclusion

Generative AI is at a critical stage. The United States must take advantage of the opportunities it presents and ensure that our adversaries don’t get a jump on us. As the applications of generative AI reach the masses in profoundly new ways, government agencies should be focused on accelerating technologies, not inhibiting them, while also ensuring that we protect our flank.

The Forum for American Leadership (FAL) is a non-profit organization that presents expert analysis and national security recommendations to policymakers in Congress and the Executive Branch.



