Written Statement for AI Insight Forum: National Security

Alexander C. Karp
CEO and Co-Founder, Palantir Technologies Inc.
December 6, 2023

Introduction

Senator Schumer, Senator Rounds, Senator Young, and Senator Heinrich, Honorable Members of the United States Senate, thank you for the opportunity to join this distinguished forum to discuss what will be America's most pressing opportunity and challenge over the next decade: How to harness the incredible power of artificial intelligence (AI), responsibly, boldly, to the benefit of our society, and in defense of our nation.

This is both a noble and herculean task. I commend Senator Schumer for his leadership in creating the SAFE Innovation Framework, as well as the corresponding AI Insight Forum series, of which I was a proud participant in the first meeting this past September. Today's meeting — on AI in the domain of U.S. national security and defense — is an important extension of our previous dialogue and I am honored to be invited to share my perspective once again.

We are here today because we understand three truths: first, that AI has the capacity to positively and exponentially improve, catalyze, and even revolutionize how individuals, businesses, organizations, and states achieve their most critical objectives. Second, that without the proper legislative, ethical, and regulatory standards for accountable and responsible AI development and use, there is a risk that AI's greatest benefits to society will be outweighed by its inherent dangers. And third, that if the United States (and its allies) do not lead in the development and governance of AI, our adversaries will, which will further put U.S. citizens and our national security at risk.

As daunting as the task before us seems, it is our duty to press ahead. The age of AI — with all its imposing benefits and risks to humanity — has already arrived, and there is no turning back. We cannot pause our efforts, both because it would be unwise in the face of adversaries who do not share our reservations, and because it would be impossible. Our adversaries abroad, which represent the greatest ideological and military threats to U.S. interests, are relentless in their pursuit to dominate AI innovation and set the global norms for its use — a reality made more troubling once we consider that in FY24, the Department of Defense will be spending only a fraction of a percent of its total budget on AI.

What is critical, however, is that our approach to advancing and harnessing the power of AI for national security should not be reactionary, or set by the purpose and pace of those that threaten U.S. interests — independent of the behavior of America's adversaries, we must recognize that AI can generate massive advancements in the readiness, resilience, and effectiveness of our Armed Forces, and thus we must seize America's first-mover advantage in this technological domain.

This is why we are gathered here today — to ensure that government, industry, and civil society work together in our common mission to ensure that AI in the national security domain is advanced safely, responsibly, and importantly, with expediency and maximum effectiveness.

Palantir has long wrestled with the realities we are discussing today. Over the past twenty years, we have securely built and deployed AI-enabled solutions — including the software architecture that makes AI effective and useful — across a diverse range of enterprises in industry and government. From public health to transportation, from scientific research to fraud detection, and from humanitarian aid to

renewable energy, our engineers are on the front-lines of society's most critical missions.

One area in which we are particularly experienced is in the provision of emerging technology to support the U.S. national security community in deterring and defending against both existing and emerging threats. I therefore feel equipped to speak directly on how the U.S. Government — along with its partners in industry — can responsibly leverage AI systems to bolster American defense.

General Principles for Responsible AI

Before offering specific recommendations on how Congress can strengthen the procurement, development, and implementation of AI systems within the domain of defense, it is important to highlight the following general principles for responsible, ethical, and effective AI development and regulation.

1. **Prioritize people, place humanity first.** The north star for responsible AI innovation is that systems should be designed and deployed to better humanity, and — to the greatest extent possible — empower, not replace, human decision-making; and that AI should be used to explicitly benefit the equitable prosperity of individuals, the environment, and humanity at large. AI should empower workers, help safeguard consumers, and enhance national security. Ethical AI is also more effective AI. The day we forget *why* we innovate — to benefit humanity — is the day the gains from AI succumb to its dangers.
2. **Operational AI — that actually works in real-world conditions — needs a robust digital infrastructure.** To build effective, safe, secure, and responsible AI, it is critical to recognize that AI systems are only as good as the broader digital foundations upon which they are built. One cannot simply invest in building and deploying models alone — our twenty years of experience in this domain has taught us that the key to operational and accountable AI is that it is built upon a robust digital architecture which includes everything from end-to-end data management to integrated hardware capabilities, as well as a broader ecosystem of established and emerging technologies. AI alone decays rapidly. In order to take meaningful steps toward enabling safe and reliable AI for American consumers, governance, and the defense community, we are best served by addressing the technologies and best practices that go into digitizing data, integrating data sources to create interoperable ontologies, implementing governance and security models, and making broader IT investments.
3. **Context matters, sector-specific regulation is essential.** The most meaningful AI risks and benefits emerge in the context of their specific use cases. Different industries, sectors, and communities each carry their own unique histories, norms, rules, resources, biases, best practices, and challenges — as such, no single standard can credibly account for all contexts, at all times. While high-level, blanket regulatory standards may be a necessary component of responsible AI and accountability, they will not be nearly sufficient to address the challenges in front of us. Instead, we must embrace the diversity of opportunities and challenges that each sector faces, and do our best to ensure that those who are closest to the ground truths of their own industries are empowered with the tools and resources *they* need to create the most effective, reasonable, and ethical regulatory standards possible.

AI for National Security: Opportunities & Recommendations

In the domain of national security, I have long believed that American innovation and software — enabled by AI in its various forms — is *the* number one advantage the U.S. has over its near-peer competitors. AI and AI-enabled software is what permits our intel officers, military personnel, and political leaders to make highly informed, data-driven decisions at a speed and scale that far surpasses our adversaries. Our unparalleled ability to maintain exquisite data-based situational awareness, to collaborate

securely across great distances, and to make data-informed decisions as events unfold, is what allows the U.S. to learn, adapt, and act at an unmatched pace. This is our ***software advantage***.

AI and AI-enabled software will not only facilitate faster and smarter decision-making, it can also significantly **magnify America's existing military advantage** by improving interoperability between hardware systems, streamlining dynamic operational planning between global mission partners, and bolstering supply-chain resilience.

Beyond these general benefits, there are countless ways in which AI and AI-enabled capabilities can improve the operational effectiveness and security of our Armed Forces. To provide just a few examples:

- AI and AI-enabled software can maximize the safety, security, efficiency, and precision of legacy defense systems, as well as underpin the development and use of future platforms and weapons systems.
- AI and AI-enabled software can improve joint- and multi-national operations and intelligence sharing, for example, by enabling real-time translation between partners using different languages, as well as by expediting classification reviews so data can be shared faster.
- AI and AI-enabled software can facilitate the automated detection of adversarial movements, including unusual ship activity and ground force deployments.
- AI and AI-enabled software can maximize force readiness and durability by improving predictive maintenance, streamlining supply chains, and minimizing budgetary waste.
- Large Language Models (LLMs) can lower the barrier to entry for the defense and intelligence community to use advanced software by allowing users to conduct analysis and give computational commands through natural language (i.e. not code). LLMs thus allow more government personnel to incorporate advanced AI-enabled models in their workflows without the need to invest in and rely on the upgraded technological proficiency of each end-user.
- AI and AI-enabled software can help strengthen an adherence to Just War and International Humanitarian Law (IHL) principles by improving the speed, clarity, and accuracy of battlefield situational awareness — potentially limiting civilian harm — as well as improving the strength of post-engagement investigations into potential IHL violations.

AI thus has the ability to take America's software advantage in defense to a whole new scale and speed, maximizing our ability to deter foreign threats, and if necessary, defeat them on the battlefield.

However, while our adversaries — including China and Russia — recognize the transformative military benefits of AI, and are rapidly integrating AI into their military operations, **the U.S. is still *not* taking full advantage of its technological capabilities.** Although the trend is improving, the Department of Defense should commit more resources to the acquisition and integration of AI and AI-enabled capabilities. For example, while the President's FY 2024 budget request for the DoD stands at $842 billion, only $1.8 billion is designated for AI, or about 0.2% of the total.

This is a dangerous reality. If the U.S. does not rapidly prioritize the sustained development and deployment of America's technological advantages for defense — primarily in software and AI — the U.S. runs the risk of: (1) Falling behind its near-peer competitors; (2) Experiencing a decline in the potency of U.S. and allied deterrence; and (3) Losing the next great power war should deterrence fail.

So how can the U.S. maximize the national security benefits of AI while safeguarding against its inherent risks?

First, the guiding principle that AI should empower humans and protect humanity is particularly true in highly consequential environments, such as national security and defense, where lives are on the line and

where culpability and accountability are the core tenets of responsible action. As such, when considering AI development and use in the national security domain, it is our collective responsibility to reaffirm — *not* release ourselves from — the need for human oversight to uphold democratic norms of privacy, civil liberties, and personal safety.

Second, we must increase our investment in the development, testing, and use of AI systems in defense. Simply increasing the amount of the FY24 DoD budget set aside for AI to just 1% (or $8.42 billion) to support our troops with the most advanced forms of commercial software available will have an outsize impact on our defense and deterrence capabilities.

Third, it is essential to recognize that, like AI use in all other domains, the only way to create operationally effective and responsible AI in the realm of defense is to continuously expose models (and their supporting software infrastructure) to realistic scenarios for testing and updating. Real-world testing, evaluation, and innovating are the only way to tailor models to perform under the conditions in which they will be expected — during live decision-making, across warfighting functions, and in real-world operations.

Funding opportunities for responsibly-constructed, "field-to-learn" experiments is an effective way to expose technologists, ethicists, policy-makers, and AI users to the specific challenges of AI deployment and use in the defense domain. For example, the recent revival of the the Global Information Dominance Experiments (GIDE) by the DoD's Chief Digital & Artificial Intelligence Office has been successful in this regard, demonstrating that the operationalization of Joint All-Domain Command and Control (JADC2) is within commercial procurement and employment reach. Toward this end, Congress can provide the Joint Staff and Combatant Commands (COCOMs) with further resources and authorities to procure and tailor commercial AI technologies to their unique geographic and/or functional needs.

Fourth, given the current comparative advantage of the commercial sector in the domain of AI — as well as the DoD's need to rapidly develop, test, and field AI capabilities — it is only prudent that DoD give due consideration to procure commercial solutions when they are readily available for use at less cost and at greater speed than a GOTS solution. As such, building on Federal Acquisition Regulation (FAR) Part 12 and the Federal Acquisition Streamlining Act (FASA), 10 USC § 3453, Congress can mandate that for any acquisition in which AI software is the core capability *and* a commercial solution is available, the DoD should be obligated to procure the AI solution from a commercial provider.

Further, Congress can mandate that software solutions be procured from providers whose core expertise is software. For platforms and systems that are defined by their software, commercial software providers usually have the capacity to stand up a minimum viable solution at the beginning of a program for immediate fielding, testing, fixing, and updating. Software firms are most likely to have the deep expertise required to play a software-defined systems integrator role (i.e. "software prime" or "software integrator"), and thus are most likely to bring AI-enabled solutions to legacy weapons systems at the greatest speed and efficiency. Furthermore, selecting software firms to provide software capabilities can save the DoD (and tax payers) billions of dollars by avoiding redundant R&D that has already been achieved by private sector investors.

Fifth, we must make multi-generational investments in the next generation of commercial defense contracting firms that are on the cutting edge of the newest technologies. Younger, non-traditional tech start-ups still face incredible barriers to entry into the defense ecosystem, and it is our collective responsibility — both the U.S. Government and established defense firms — to help these innovative newcomers bring their solutions to the U.S. Government's most pressing problems. We need to support the health of America's defense tech ecosystem**.**

Beyond the standard calls for procurement and budgetary reform — which Congress is already heeding through its PPBE Reform Commission process — it is important to address a less visible, yet equally constraining feature of the "valley of death" for young firms: The length, costs, and complexity associated with the accreditation and compliance process for new technology. In its current state, this process hinders the DoD's capacity to rapidly absorb innovation like AI by making it excessively difficult for start-ups to pursue and achieve Impact Level accreditation and FedRAMP access in the first place. One way to potentially resolve this innovation procurement choke-point is to outsource the most time-consuming and costly tasks associated with this process to established commercial firms, who can facilitate the onboarding of new firms' technologies into the federal government domain through their existing authorizations.

Conclusion

Countries like China are investing heavily in AI and are eager to leverage its power to threaten the U.S. and its allies abroad.

We do not have the luxury to debate without action. Nor do we have the luxury to endeavor towards our AI-enabled future through independent efforts — government, industry, civil society, and academia must work hand-in-hand to not only develop the AI systems that will drive economic progress and strengthen American defense, but also to create a set of reasonable standards and regulations to ensure that ongoing AI development and use is ethical, responsible, and advances American values.

To my fellow tech leaders who remain wary of providing your capabilities and expertise to our government, I strongly urge you to reconsider. Even as industrialists, we do not have the luxury to ignore international politics. We owe our success and privilege to this country — America's founder-driven model of capitalism, its free-market environment, its unabashedly meritocratic corporate culture, and access to large pools of risk-tolerant capital and high-skilled labor has catalyzed our comparative advantage in innovation. It is this country, and its values, that is the very reason we are all sitting at this table today. The U.S. defense community is looking for a partner. In particular, they are looking for America's tech innovators, who are driving the very technological change that commands our attention.

This forum is ultimately convening in agreement on its most basic principles — that America must envision and endeavor towards an increasingly AI-enabled future. A future that is safe for consumers. A future that brings growth and equitable outcomes to all members of society. A future that safeguards U.S. national interests and individual well-being. A future that upholds our founding principles as a democratic country.

While our perspectives and recommendations may diverge in other areas, our shared understanding of why we are here and why this moment matters gives me confidence in the success of its outcome.

Thank you, and I look forward to our discussion.