AI an Essential National Security Asset: Prudent Oversight that Advances Innovation

Presented to the AI Insight Forum on National Security
December 6, 2023
Dr. Charles F. McMillan
Director Emeritus Los Alamos National Laboratory

Section 1: Introduction

Today, technological prowess increasingly determines geopolitical influence. Artificial intelligence (AI) is poised to become a new frontier in national security and in the exercise of that influence. The imperative to harness AI's potential for national security is more than a strategic advantage; it is essential. As global powers vie for technological supremacy, the United States faces a critical challenge: to lead in AI or risk falling behind in a domain crucial to future conflicts and international competition. This urgency is underscored by the rapid pace of AI development globally, and its potential applications in various domains from fundamental science to intelligence. The U.S. must not only innovate but also adapt to the evolving landscape of AI to maintain its position on the international stage.

Section 2: Governmental Oversight in Emerging Technologies

The rapid evolution of AI, marked by its expansive potential and multifaceted applications, poses a significant challenge for governance. Drawing parallels with historical precedents set during and after the Manhattan Project 80 years ago reveals both similarities and important distinctions. The early nuclear era was predominantly shaped by the government-led Manhattan Project and was characterized by high levels of classification. This era culminated in the creation of the Atomic Energy Commission (AEC) in 1946, which began nearly a decade of continued government monopoly in both the national security and nuclear energy sectors.

Based on my personal conversations with laboratory workers who began their careers in this period, it's apparent that government regulation initially adopted a 'light touch' approach. This was a time when the technology was still in its nascent stages. While government oversight was recognized as essential, the comprehensive regulatory framework that later developed was not yet in place. The government's initial strategy centered on oversight through entities like the AEC and the Joint Committee on Atomic Energy (JCAE). These mechanisms struck a balance between the need for flexibility and the necessity of innovation in a rapidly evolving field.

In stark contrast, the emergence of AI has predominantly occurred within the commercial sector, and is marked by near-exponential growth. This divergence requires a different approach to governmental oversight compared to the early nuclear era. However, there are valuable lessons to be learned from those formative years that could be relevant in today's AI landscape. Key among these is the importance of ensuring that societal interests are protected while simultaneously providing the regulatory leeway necessary for the growth and maturation of this emerging technology.

In the global arena, AI development is both a cooperative and competitive endeavor. AI advancements are taking place on a global scale, involving a multitude of players from different nations. This presents both opportunities and challenges for governance.

On one hand, international cooperation is essential in setting standards and regulations for AI, especially when considering its implications for global security, ethics, and economic impact. Cooperative efforts, such as international treaties and agreements, could play a pivotal role in establishing a common ground for the use and oversight of AI technologies. However, should treaties be beyond reach in this time of geopolitical tension, institutions such as the International Atomic Energy Agency (IAEA) that pre-dated all of our major arms control agreements, might still be useful. This would not only help in addressing shared challenges like cybersecurity threats but also in ensuring that AI development progresses in a manner that is beneficial to the international community.

On the other hand, the competitive aspect of AI cannot be overlooked. Nations are vying to establish supremacy in this rapidly evolving field, viewing AI as a key component of their geopolitical and economic strategy. This competition drives innovation but also raises concerns about a potential arms race in AI technologies, reminiscent of the nuclear arms race of the last century. Effective governance in this context must therefore navigate the delicate balance between fostering healthy competition and preventing unchecked AI advancements that could lead to global instability.

Public-private partnerships can serve as a catalyst for innovation in AI while ensuring that development aligns with national security interests and ethical standards. Such collaborations can provide a framework for sharing resources, knowledge, and expertise between the government and private entities. This synergy can accelerate advancements in AI, from research and development to practical applications in national security.

Moreover, these partnerships can facilitate a more nuanced approach to regulation. By involving stakeholders from the industry in the governance process, regulations can be crafted that are both effective in mitigating risks and flexible enough to accommodate the fast-paced evolution of AI technologies. These partnerships also open the door for leveraging private sector innovation to solve public sector challenges, creating a win-win situation that advances national interests while promoting a thriving AI industry.

Section 3: Scientific Applications and Operational Enhancements in National Security

AI's Emerging Role in National Security

AI's applications in national security, though still in their developmental stages, are showing significant promise. Its potential ranges from enhancing data analysis for improved decision-making to streamlining manufacturing processes. The following areas are particularly promising for the application of AI technology:

- **Surveillance Data Analysis in the Nuclear Stockpile**: AI's capacity to process and analyze vast datasets can play a crucial role in nuclear security. By identifying patterns that indicate concerns within the stockpile, AI can provide early warnings, long before issues become apparent to engineers. This proactive approach to surveillance could be pivotal in maintaining the safety and integrity of our nuclear arsenal.
- **Manufacturing of Nuclear Components**: The manufacturing process for nuclear components, especially pits, generates extensive data. AI's ability to analyze these data may be able to detect early signs of manufacturing anomalies, allowing for prompt corrections and ensuring smoother

operation of manufacturing processes. This application not only enhances efficiency but also contributes to the reliability of critical national security assets.

- **Assessment of Chem/Bio threats**: These are areas where the National Security Laboratories have partnered with the Department of Homeland Security (DHS) and have applied their computing prowess not only to national security issues, but also in support of management of the Covid epidemic. AI presents the possibility of applications like predictive epidemiology, something that is important for both national security and public health.
- **Rapid Code Development with AI-derived Large Language Models (LLMs)**: LLMs have demonstrated their ability to write high-quality computer programs, mirroring the leap in productivity seen when scientists moved from machine language to high-level compilers. AI tools now offer a higher level of abstraction, enabling scientists to write code more efficiently and focus on complex problems rather than the intricacies of programming languages.
- **Code Optimization for Scientific Algorithms**: Recent research[1] has shown that AI techniques can significantly enhance the performance of even the most optimized algorithms, such as sorting algorithms. This opens possibilities for similar improvements in scientific algorithms that are used in national security, particularly in simulations related to the performance of the nuclear stockpile. These algorithms have not been optimized as extensively as sorting and may have even more potential for improvement. Such advancements could lead to more accurate and efficient simulations, providing critical insights for national defense.

These examples highlight the potential contributions of AI and ML technologies in the national security laboratories. However, it's important to note that, like the Manhattan Project, the data sets and results from this work will be highly classified and distinct from the commercial sector's large-scale AI training platforms.

Linking Past and Future: From Game Theory to AI in Deterrence

The work of deterrence theorists like Thomas Schelling and Herman Kahn in the 1950s and 1960s was informed by their work on two-person games of perfect information. This approach was apt for an era dominated by two superpowers. However, the current geopolitical landscape, with multiple nuclear powers, requires a shift to multipolar theories of deterrence.

While AI's proficiency in games of perfect information like Chess and Go is well-documented, these simulations involve only two players and thus have limited applicability in our increasingly multipolar world. In contrast, recent advances in AI-powered poker[2], a game of imperfect information with multiple players, are particularly intriguing. These algorithms have achieved proficiency, routinely outperforming human players, demonstrating AI's potential in more complex, multiplayer strategic scenarios.

The question arises: could we design multiplayer games that mimic the intricacies of international deterrence? Could insights from these AI-powered games inform and advance our deterrence theories beyond the foundational work of Schelling and Kahn? While the answer remains uncertain, the likelihood that our competitors might explore this avenue suggests the urgency for us to engage in such

[1] Mankowitz, D.J., Michi, A., Zhernov, A. *et al.* Faster sorting algorithms discovered using deep reinforcement learning. *Nature* **618**, 257–263 (2023). https://doi.org/10.1038/s41586-023-06004-9

[2] Brown, Noam, and Tuomas Sandholm. "Superhuman AI for multiplayer poker." *Science* **365**, no. 6456 (2019): 885-890.

research. This proactive approach could be crucial in shaping future strategies for a world with diverse nuclear capabilities.

Section 4: Path Forward: Drawing Parallels and Learning from Historical Models

The lessons learned during the early nuclear era offer valuable insights for contemporary AI governance. Historical models like the Atomic Energy Commission (AEC) and the Joint Committee on Atomic Energy (JCAE) illustrate how transformative technologies can be overseen responsibly, especially during periods of rapid evolution. Although these models were designed for a different era and technology, they underscore the importance of adaptable governance structures capable of evolving alongside the technologies they regulate.

In the context of AI governance, the role of the National Nuclear Security Administration (NNSA) becomes particularly significant. Given the NNSA's responsibilities in ensuring the safety, security, and reliability of the nation’s nuclear deterrent, coupled with the expertise of the National Security Laboratories in computing, it is crucial for these entities to be at the forefront of AI application development in national security. The current management structures available to the Department of Energy (DOE) and NNSA, particularly the model of Federally Funded Research and Development Centers (FFRDCs), offer a blueprint for this endeavor. These labs, operating as Government-Owned, Contractor-Operated facilities, represent a fusion of public accountability and private sector flexibility and innovation. This hybrid model can serve as an effective framework for AI oversight.

Adapting these structures for AI governance would entail prioritizing flexibility to accommodate the rapidly changing landscape of AI technology. It would also involve fostering public-private partnerships, leveraging the innovation of the private sector while ensuring that development aligns with national security objectives. Such a governance model would not only facilitate the advancement of AI in critical areas but also ensure adherence to ethical standards, balancing technological progress with societal and security concerns.

This approach, drawing on historical and modern governance structures, can guide the development of a robust AI governance model that effectively navigates the unique challenges of AI in the national security context.

Section 5: Challenges and Resource Requirements for Integrating AI in National Security

The successful integration of AI into national security entails overcoming several significant challenges. Foremost among these is the need for adequate funding. Investment is crucial not only for research and development but also for establishing the necessary infrastructure to support AI applications. This includes the creation of secure and robust computing resources capable of managing the immense data requirements intrinsic to AI technologies. Current funding programs did not anticipate the swift advancements in AI and Machine Learning (ML) that have emerged in recent years. While AI is expected to augment the ongoing work under the program of record, it is impractical to rely on reallocating funds from vital projects aimed at modernizing the nuclear stockpile. Dedicated funding sources will be necessary to support AI initiatives without compromising essential ongoing projects.

Workforce recruitment and development is another pivotal area. The national security sector requires a cadre of professionals who are not only adept in AI and related technologies but also fully comprehend

the unique requirements of national security work. This necessitates a concerted effort in education and training, bolstered by partnerships between government entities, academic institutions, and industry. Given the high demand for AI expertise in the commercial sector, attracting top talent to national security roles will require leveraging the flexibility afforded by Government-Owned, Contractor-Operated (GOCO) institutions. It will be necessary to offer competitive incentives and career development opportunities that align with the critical nature of this work.

Furthermore, the dynamic nature of AI advancement demands that the regulatory and oversight frameworks be both agile and visionary. They must be capable of adapting to rapid technological changes and unforeseen developments, while simultaneously ensuring security and ethical standards are upheld. This necessitates a continuous dialogue among a diverse group of stakeholders, including technologists, policymakers, security experts, and ethicists. Such collaborations are essential to foresee, understand, and effectively navigate the emerging challenges and opportunities presented by the integration of AI in national security.

Section 6: Conclusion

As AI becomes increasingly integral to national security, it is imperative that the United States positions itself as a leader in this domain. This requires not only continued investment and innovation but also prudent oversight that fosters growth while ensuring safety, security, and ethical integrity.

We need policies that support robust AI development for national security. These include explicit funding for AI research with a focus on national security applications, incentives for public-private partnerships, and the establishment of a national framework for AI education and workforce development.

There may be value in considering structures, such as the Joint Committee on Atomic Energy that were both bicameral and bipartisan, which have served us in the past and helped us to navigate the unique challenges of new technologies while promoting innovation. This body could be responsible for developing flexible regulatory frameworks, fostering international cooperation in AI, and ensuring that AI development aligns with national security objectives and ethical standards.

In conclusion, the United States stands at a critical juncture. The decisions made today will shape the role of AI in national security for years to come. It is essential that these decisions promote innovation while ensuring responsible oversight, securing the nation's position as a leader in this pivotal technological domain.