SPECIAL COMPETITIVE STUDIES PROJECT

# U.S. Senate "AI Insight Forum" on National Security

*Written Statement by Dr. Eric Schmidt*

*"The AI-Enabled Future of U.S. National Security"*

Leader Schumer, Senators Rounds, Heinrich, and Young, and distinguished members of the Senate, thank you for inviting me here today.

I had the honor of attending the first AI Insight Forum, where I shared my thoughts on how Congress could engage in AI policy by promoting U.S.-led innovation and exploring appropriate safeguards. I'm here today to provide insight into the latest developments regarding AI — notably, large language models — and their impact on national security.

As many of you know, I founded the Special Competitive Studies Project, a nonpartisan, nonprofit that makes recommendations to strengthen America's long-term competitiveness as AI and other emerging technologies are reshaping our national security, economy, and society. My vision for our nation's future has been shaped by my time leading the National Security Commission on Artificial Intelligence (NSCAI) and the Defense Innovation Board (DIB), as well as through my philanthropic work and time in the private sector. Throughout my experiences, I've come to realize the future of AI is more exciting, and the threats are more concerning than I ever thought possible.

The AI industry is undergoing an investment and growth cycle on a scale we have never seen. The systems are scaling quickly in capacity and impact. The potential for universal AI applications in healthcare, education, science, and national defense is immense; however, the same advancements pose significant challenges in terms of proliferation and control, especially when AI becomes a tool in the hands of those who might wish our nation, or our world harm. The rapid evolution of AI technologies, including generative design, is revolutionizing not just industry but the very nature of warfare and security. The unpredictable trajectory of AI development, compounded by reduced costs and increased accessibility, raises critical questions about proliferation, governance, and safety.

Today, we have widespread access to many large language models, with many more coming soon. They are generally divided into Frontier models, which are large and very expensive to train (like GPT-4 from OpenAI, Claude2 from Anthropic, Gemini from Alphabet, and PI from Inflection), and smaller models, which I call mid-tier models, that are also very powerful but less expensive and often open sourced (including LLAMA2 from Meta, 7B from Mistral, and many other companies in which I am an investor). In the very short term, we can expect these models to mature quickly, and many of their current issues, such as hallucination, will be addressed. The consensus today is that the mid-tier models lag the frontier models in functionality by

about eighteen months to two years, which means the hottest new discovery in one company is broadly available globally within a few years.

The power of models will increase 1,000 to 10,000 times over the next decade. With ten times more investment, more powerful hardware, and significantly better algorithms, the changes and improvements will be exponential and largely unpredictable. These systems will be polymathic, learning, imprecise, and unpredictable.

Together, we must navigate these uncharted waters with optimism for the potential benefits and vigilance for the unprecedented national security challenges ahead. New models and new actors will emerge from countries and groups because they understand the importance of these models and the transformative power they can have for their nations. Our nation's strength and future prosperity hinge on our ability to stand at the forefront of science and technology development and ignite innovation power; therefore, as we proposed at the NSCAI, we must make a significant investment of $32 billion in non-defense AI research and development to match Cold War spending as a percent of U.S. gross domestic product.

**The Risks and Benefits of AI in National Security**

*The Proliferation Problem*

The widespread availability of inexpensive and powerful models is a large-scale proliferation problem. Today's models are tested for AI safety after the training is completed, and guardrails are then added during this final phase, often by human experts. These guardrails appear to be easy to remove by a sophisticated attacker, and we should expect the availability of powerful models that may be a platform for misuse, at least on the dark web. The top alignment training sets will likely be mostly open, as we benefit from seeing what tests they are subjected to, and others can help find errors and omissions in the training and testing regime. However, the issue is that some dangers are serious enough that the tests themselves should not be published. Ultimately, we must establish some thresholds of capabilities for the future – if the model can do X, then we should stop and consider how to handle this new capability.

The competition between so-called closed frontier models and smaller open source models will be fierce. Proponents of the closed model argue that their superior resources, including hardware, software, and scale, will enable them to align more closely with human values and surpass open models in capability within the next two to three years. The open model proponents claim that their inherent transparency makes them more suitable for government and other uses, and they claim they will move faster and be more innovative than their larger closed models. They argue that they can coalesce and combine open models and compete with greater safety and power than the closed models. This will be a contest for the ages — both models are being funded at aggressive levels, and huge capital bets are being made. The

result will be a very, very large number of models of different sizes, each of which has the potential for misuse and uncertain impact.

*Generative Design Revolution*

The newest systems are multi-modal and develop an ability to analyze and generate new images, facts, fallacies, and history. Generative design, where the computer generates something similar to what humans have done, will change art, music, entertainment, storytelling, politics, science, and engineering. But many businesses will also use it to increase business efficiencies, as the personal computer did 50 years ago. My favorite example is the insurance business that rejects an insurance claim with an AI-generated letter. The AI responds with its appeal counterclaim letter, and then the AI system decides if the appeal is to be accepted. This is progress, I think. What happens to all of those jobs? The great artists will continue, but what about the humans who more or less repeat what others have done? What about the humans who work in large accounting and auditing functions, in internal processes that can be replaced? When their jobs are replaced, where do they go? My instinct is that AI will eventually double human productivity — if you have a job. What will happen when our nation's productivity doubles? What happens if our adversaries' productivity doubles?

We expect sequential planning to be possible in the next one or two generations of LLM platforms. The system will be able to make predictions that are consistent with one another. Through an application programming interface (API) interface, the agent representing me can become very powerful. We can expect future systems to have the ability of a senior project manager. These models replace labor for many tasks but are also certainly capable of being regulated. We will simply apply the same laws we have for humans for their agents. We may also be able to build agents that are limited in what they can do if the actions are clearly evil or illegal. With sequential planning, we can have a huge job impact, but we can also hold the owners of the systems accountable.

*The Future of War and AI*

War has been constant in human history, and AI will not solve this. However, during my trips to Ukraine, I witnessed firsthand how technology — including AI — is making a difference on the battlefield. The war in Ukraine is a technical war fought by technical people. And war is now fought at a much more rapid pace on a highly decentralized battlefield, with new and old capabilities employed in incredibly innovative ways. The combination of proliferated sensors, autonomy, vast improvements in drone technology, and the use of commercial space and software are changing war as we know it. There is no place to hide anymore. If you co-locate, you get destroyed. If you congregate, you die. This reality demands we rethink how we design, organize, train, and equip our military. We are seeing an emerging combination in Ukraine of rapid and accurate target recognition and algorithmic management of battlespace awareness and battlefield assets that is constantly optimizing assets to win — just as we saw AlphaGo

complete in a game. Proliferation will be a huge issue here, and so will a new arms race where each side improves its offensive army while improving its defenses against the same.

A $5,000 drone can easily destroy a $5 million tank and should be able to replace most tanks and even artillery in war. Land wars and invasions between AI-armed peers will become much harder to win as the other side will attack the invading forces immediately. Countries will still have traditional militaries, especially ground forces, but they will follow after the robotic war and, depending on who is winning, may or may not be able to win on the ground. A strategy of networked war allows drones to be nearly invincible, and a strategy of abundance and software reliance means there will be many ways to win.

It's unlikely that governments will agree on the restrictions for robotic war, but we can probably get an agreement to require meaningful human control over these actions, as we saw from the meeting between President Biden and President Xi Jinping in California earlier this month. The core principle of the laws of war — that people can be held responsible — must be preserved with the new arrival of AI, drones, and algorithms in warfare.

**A Way Forward**

The stakes for the United States have never been higher. Our armed forces' competitive military-technical advantage — which we have enjoyed for over 70 years now — could be lost within the next decade if we do not accelerate the adoption of AI applications across their missions and develop new operating concepts to exploit the power of AI. I firmly believe that the technological capabilities in the private sector should be available to our military. To this end, the Department of Defense (DoD) has taken important steps to adopt private-sector technologies. It has established institutional structures to mainstream AI across the Pentagon. It is beginning to pursue low-cost drones at scale. And it has promulgated policy guardrails for human-machine teaming in warfare. The Department now works more closely than ever with new companies, including small-scale startups previously outside the defense sector. These efforts increase competition and creativity and create a more agile, flexible Department capable of addressing the strategic challenges ahead.

Many of these initiatives have their genesis in the work of the DIB. Here, we recommended a number of steps that the DoD should pursue to generate military advantage against advanced peer adversaries. The DIB's Software Acquisition & Practices report was instrumental in launching the "software revolution" across DoD. DIB's recommendations helped launch Project Maven, a pathfinder project for subsequent AI initiatives. The DIB's work also paved the way for the first AI Strategy and the creation of the Joint Artificial Intelligence Center (JAIC) at the Pentagon, which has now evolved into the Chief Data and Artificial Intelligence Office, chartered with accelerating the Pentagon's adoption of data, analytics, and AI to generate

decision advantage. The DIB also looked beyond AI, exploring risks and opportunities for DoD on 5G, cyber, and talent management, to name a few.

In the NSCAI Final Report, we recommended that DoD be fully AI-ready by 2025. Becoming AI-ready requires the DoD to invest in R&D, but it also requires engaged leaders who shepherd the development of innovative operational concepts and business practices that take advantage of AI's speed, scale, and optimization. To drive such adoption, the DoD needs to establish AI-readiness performance goals focused on logistics, experimentation, and use in training. Just as importantly, the DoD needs leadership empowered to hold accountable organizations that do not meet their performance goals.

Today, driven by the fundamental changes in the character of war we are witnessing in Ukraine, along with China's accelerated military modernization and hardened ambitions, we must lay the groundwork to maintain or regain military-technological superiority. That means developing entirely new operating methods enabled by advanced technologies — specifically AI.

At SCSP, we have articulated a new competitive strategy for DoD to help us achieve these goals, called Offset-X, that recommends our military operate as a distributed, network-based force; leads the world's militaries in human-machine collaboration and combat teaming (e.g., networked drones assisting every human soldier); and develops, deploys, and updates innovative software that will help us defeat any potential adversary.

Just as AI has the power to transform our military, it too has the power to transform the business of intelligence. A revolution in open source intelligence powered by big data analytics at speed and prioritization of techno-economic intelligence will fuel information advantage for our policymakers, which will be much needed as competition with the PRC increases. These technologies will enhance our capabilities, but they will also amplify threats from our adversaries through AI-powered malign disinformation, enhanced cyber capabilities, and more.

As we stand on the precipice of an AI-dominated era, we must prepare our people, systems, and nation for both the challenges and opportunities ahead. Our approach must be a holistic whole-of-nation effort encompassing not just technological advancement, but also the training and development of our defense and civilian workforce to navigate and leverage AI adeptly. Equally important is maintaining an ethical and responsible approach to AI deployment, upholding our democratic values, and pursuing truth, even as we confront these new digital battlefields.

Thank you again for the opportunity to appear before you. I look forward to our discussion.

# # #