Written testimony of Dr. Dewey Murdick
Executive Director
Center for Security and Emerging Technology, Georgetown University

For a Senate Homeland Security and Governmental Affairs Subcommittee on Emerging Threats and Spending Oversight hearing on Advanced Technology: Examining Threats to National Security

September 19, 2023

Chairwoman Hassan, Ranking Member Romney and honorable Senators of the Emerging Threats and Spending Oversight Subcommittee, thank you for the opportunity to discuss the increasingly vital topic of how emerging advanced technologies are affecting our national security. Many of the ideas I will discuss are motivated by the data-driven, tech-policy analysis from Georgetown's Center for Security and Emerging Technology. Others come from my own experience working within government departments and agencies, a couple years living and working in Silicon Valley, and academic experiences.

Key Questions About Emerging Technology Threats

Elected officials and public servants are continuously bombarded with warnings about looming threats or game-changing technologies, all demanding urgent action and investment. Many of these warnings and promises are based on something real, but how do we decide which are most relevant and deserve the most attention and resources? There are three important questions that I think all policymakers should ask when considering various emerging technologies, proposals to address their threats, and recommendations for capitalizing on their potential. These questions can help prioritize our attention and national resources toward the most urgent and transformational efforts and include the following:

1. **What technologies are most timely (i.e., exist now or will exist soon) and have the most significant impact on our national security, economic competitiveness, and societal well-being?**

   Many of today's emerging technologies require urgent action, a demand that this body is rightfully working to meet. For example, AI could improve a terrorist's target reconnaissance and attack planning, allow an adversary to generate more and better disinformation or phishing attacks, guide criminals on the best way to acquire chemicals for meth production without triggering law enforcement alerts, or assist human traffickers in developing more convincing pitches to manipulate victims or their families.

   Others, such as the threats from superintelligent AI systems, are more nebulous, often because we don't yet understand enough about the potential or limitations of the technology involved. I am no

expert in quantum computing implementation, but CSET staff have talked with experts and concluded that the practical ability to break through protections and read messages encrypted by modern algorithms, a key application area, remains potentially as far as decades away. As of now, existing quantum computers have nowhere near the scale and capacity required to tackle these challenges.

And still others may not be as transformative as we might think. Recent coverage of how AI chatbots could potentially assist non-experts in dangerous scientific endeavors such as creating pandemics, sparked by a recent preprint article, may exaggerate the threat. An individual motivated to cause harm could already access a wealth of scientific information and protocols online just like the MIT class in this article did. They could learn basic lab techniques from YouTube tutorials, guided by the help of an animated beaver, or even order high-school level kits (in magenta no less) designed to teach bacterial engineering. Open scientific learning and discovery is a very good thing, so I am not advocating that this be curtailed; however, the information barrier for bad actors is not as high as it might appear, and chatbots at this stage serve more as a convenient tool that brings together information in context rather than a game-changing enabler.

In the case of large language models, it doesn't seem that advances here are as relevant to competition with China as we may initially think. Given the Chinese Communist Party's penchant for tight control over the domestic information environment, unpredictable generative AI that might generate responses on Tiananmen Square or images of Winnie the Pooh are unlikely to be a priority. Instead, their efforts are more inclined toward surveillance technologies, such as computer vision methods used for monitoring the face or walking style of tracked individuals.

2. **Do we have the plans, authorities, and tools necessary to mitigate the threat or capitalize on the opportunity? If not, what can we easily update?**

In the short term, there are steps we can take to promote the development of AI that aligns with U.S. values. We should find ways to get the information we need to make better decisions by:

- Tracking AI harms in incident reporting (including voluntary and required reporting);
- Examining the data and models used in existing widespread applications to increase accountability (e.g., sentencing algorithms);
- Encouraging the development of the third-party auditing and red teaming ecosystem; and
- Improving the quality and security of resources, including training and pretrained models that form the backbone of many of today's AI systems.

As we look forward into the near future, it is crucial to **take stock of our existing relevant authorities** that apply to new application areas and leverage the existing strengths of our nation. Understanding and effectively using these existing authorities and abilities will help us move

forward and determine which areas require updates or additional focus. To my knowledge, these assessments have not yet been completed.

In the longer term, additional authorities — or even a new agency — may be needed to lead coordination, develop a critical mass of expertise, or improve oversight. A new agency could:

- Check how AI is being used in and overseen by existing agencies;
- Be the first to deal with problems, directing those in need to the right solutions; and
- Fill gaps that existing sector-specific agencies don't cover.

However, Congress should be cognizant of the time and money it will require to stand up a new agency, let alone allow it to gain its political footing in the D.C. bureaucracy. This doesn't mean it shouldn't be done, but it also shouldn't hold us up from taking other action now.

In the case of biotechnology, the primary concern should not be the ease of information access facilitated by chatbots, but rather on strengthening biosecurity regulations that govern physical experimentation and tangible materials, including access to tools and resources like custom mail-order DNA. Currently, DNA synthesis providers are not required to screen orders for sequences of concern, allowing bad actors to obtain materials with minimal oversight. The 2010 Health and Human Services Screening Framework Guidance offers voluntary guidelines for DNA synthesis providers, but adoption remains optional. A more effective approach to risk mitigation would be to mandate screening protocols for all DNA synthesis providers, thereby directly addressing the foundational risk pointed out by the MIT experiment.

Additionally, we can and should develop emerging technology-aware talent. Doing so will take time, but we cannot afford to underutilize our human capital and it is this long horizon that requires us to act now. Congress can help right away by facilitating a targeted high-skilled immigration program and incentivize efforts to advance general AI literacy, grow more STEM talent, and promote certification programs within the United States. While most of these ideas are focused on AI, many of these same principles apply to cybersecurity talent and the broader biotechnology workforce.

3. **Are our plans, authorities, and tools easily adaptable if and when the landscape changes? Do we have the information necessary to know when that adaptation needs to happen?**

Finally, we need to be prepared to adapt these plans as technological breakthroughs occur or as the threat environment changes. In order to adapt to new information and update our priorities based on changing landscapes, we need a dedicated monitoring capability that tracks emerging technology developments, both here and abroad. A well-funded, decision support capability (see an early proposal) using publicly available information can tell us things like:

- Who has essential capabilities in vital research and the transformative knowledge poised to change the emerging technology landscape;

- What problems our adversaries and competitors are funding and trying to solve and what this tells us about their strategies and priorities;
- Where new technologies are being applied, such as the use of AI in genetics or other fields, where risks might be particularly concerning;
- Talent development trends, including what fields are attracting talent and where those skilled workers are going to work after they receive their education and training, both in the United States and abroad; and
- The makeup of supply chains and any associated checkpoints that jeopardize our security or that we can leverage toward others.

Then we need to be able to effectively coordinate information sharing and policy adaptations across agencies, based on the information provided above. This requires resourcing and expertise embedded in core government functions, as opposed to being only run by time-bound political appointees.

Finally, asking these questions is not a static exercise and constant reevaluation is needed. What may be lower priority today can easily jump toward the top of the list next month or in two years. Given the sensitive and often long-lasting nature of government communications, it is prudent to continue the transition to data encryption methods that cannot be easily cracked by quantum computers and ensure that the attention paid is proportional to realistic assessments of opportunity and risk.

For this and other topics not prioritized today, the technology monitoring mechanism can ensure we are tracking technological progress and limit the possibility that we are surprised by major breakthroughs.

Conclusion: Agility and Vigilance is Key

In conclusion, our approach to emerging technologies and their impact on national security must be agile, adaptive, and action-oriented. We can't afford a "set and forget" mindset; we must be prepared to continually adjust our strategies as technologies evolve and the global threat landscape shifts. Different nations and actors have unique motivations, which could influence technological developments in ways distinct from our own. We need to be vigilant in monitoring these variances and ready to prioritize and de-prioritize as things shift.

A practical way forward is adopting an iterative strategy: take small steps, evaluate their effectiveness, learn from the results, and then take the next steps. This is not just an operational recommendation but a call to invest in analytics. Better data and insights will allow us to make more informed and timely decisions, effectively allocate resources, address regulatory gaps and build a robust talent base.

By committing to a continuous learning process, we ensure that our approach remains nimble and responsive to the rapid developments in technology that we're bound to face.

Thank you.