Written Testimony of
Dave Ferris
Head of Global Public Sector, Cohere
Before the Senate Armed Services Committee, Subcommittee on Cybersecurity

March 25, 2025

Introduction

Chairman Rounds, Ranking Member Rosen, and distinguished members of the subcommittee: thank you for the opportunity to testify today. My name is Dave Ferris, and I act as Cohere's Head of Global Public Sector. Prior to this role, I served for nearly 17 years in the Canadian Armed Forces, including several years on the Joint Staff at the Pentagon, and held various roles in private industry focused around the use of technology in aerospace, defense, and intelligence.

Cohere is a leading artificial intelligence (AI) company that builds state-of-the-art foundation models and agentic systems exclusively for enterprise and government use; with a focus on **privacy, security, multilingual capability, and verifiability**. Our team includes several AI pioneers – from our CEO & co-founder Aidan Gomez, who helped invent the transformer architecture underpinning today's advanced language models; to our Director of Machine Learning, Patrick Lewis, who invented Retrieval Augmented Generation (RAG); and our VP of Research, Sara Hooker, who has pioneered efficiency and multilingual techniques in AI. We combine deep technical expertise with a practical, mission-driven mindset. We partner with allied government agencies and leading global companies such as Oracle, RBC, LG, and Fujitsu, focusing on seamless integration, deep customization, and accessible solutions that deliver immediate practical value.

Cohere strives to be a trusted partner in the national security community. We are proud to support the United States and its allies in leveraging AI to strengthen national security across all of its missions. We focus not just on inventing novel AI models in a lab, but on **operationalizing** AI – integrating it into real missions, under real-world constraints, in a manner that delivers tangible value. We understand the unique needs of the defense community and are committed to ensuring our AI solutions are robust, trustworthy, and ready for use in these environments.

In today's testimony, I'd like to highlight four topics of focus: 1) How AI can support the Department of Defense's (DoD) mission through lessons we've learned from cyber defense deployments, 2) Key Technical Considerations for AI in Defense, 3) Lessons from the DoD's deployment of technologies in the past, and 4) Steps Congress can take to accelerate responsible AI adoption in defense.

1. How AI can Support the Department of Defense's Mission

AI is an umbrella term used frequently to refer to many different technical systems with distinct capabilities - all working in a hybrid AI ecosystem. Recognizing this distinction is important when discussing these systems in context. Machine Learning (ML) systems refer to the wider category of systems that apply statistical learning techniques to learn patterns from data and make predictions or decisions - this includes things like recommendation algorithms and anomaly detection systems. Large Language Models (LLMs) refer to systems built around large-scale language data to understand, generate content, and perform tasks. While Cohere makes LLMs and agentic systems for augmentation and automation, we posit on broader AI trends, including ML systems, in this testimony.

AI is **transforming cybersecurity, intelligence, and general defense operations**, and holds enormous promise for the DoD as it confronts rapidly evolving threats. Having worked with high security cyber defense government clients on adopting AI, Cohere has gleaned several insights into how AI can best augment the DoD's mission.

One clear lesson is that **AI can dramatically improve pattern recognition and anomaly detection across the vast datasets** analyzed by defense and intelligence agencies. In cybersecurity operations, for example, ML systems can comb through millions of network events to find the proverbial needle in a haystack – a suspicious pattern of behavior or a subtle anomaly indicating a cyber intrusion. These models can learn normal baseline activity and continuously scan for outliers, giving analysts a powerful tool to detect advanced threats (such as nation-state hackers using novel techniques) much earlier than manual methods. Similarly, LLM-driven systems can correlate disparate indicators and alert human operators to threats that would have otherwise gone unnoticed. This outcome can measurably improve threat detection, with AI providing a second set of eyes (which are faster and tireless) across complex networks.

Similarly, in the intelligence analysis realm, LLMs have proven **invaluable for sorting through and synthesizing huge volumes of multi-source information**. Intelligence analysts are often inundated with so much data – from satellite imagery to intercept transcripts – impossible for any team of humans to thoroughly examine. AI helps by

rapidly categorizing and triaging this data. For instance, computer vision models - including multimodal foundation models - can scan video feeds and flag objects or activities of interest for human review, turning hours of raw footage into a shorter list of likely threats. Language AI can automatically translate and summarize foreign communications, highlighting potential security-relevant content among thousands of messages. By improving pattern recognition and triage, AI allows human experts to focus their attention where it matters most – on the truly hard cases and strategic judgments – while machine assistants handle the initial heavy lift of data processing.

Another key insight from our work is **how multilingual and cultural breadth in AI systems can advance national security**. Threats and relevant intelligence can emerge in any language or region, yet many AI models historically have been skewed toward English and a few major languages. This creates blind spots.

We've found that adopting AI in national security contexts often requires closing this "language gap", and Cohere has developed highly multilingual language models, covering 23+ languages. Whether it's analyzing in Pashto, scanning open-source reports in Mandarin, or assisting a humanitarian mission in French or Spanish, AI needs to understand the content. Multilingual AI models dramatically expand the intelligence community's reach by enabling automated analysis of text or speech in virtually any language of interest. We've also learned that incorporating diverse languages and dialects into model training not only broadens coverage, but also improves the overall safety and accuracy of AI outputs. In practice, a multilingual model is less likely to misinterpret critical nuances, which is crucial when informing high-stakes national security decisions. Thus, the DoD and the IC should continue to prioritize AI that reflects the linguistic diversity of the real world so that our technologies can better execute any mission set, anywhere in the world.

We have also observed that successful AI adoption is not just about the technology – it **requires adapting workflows and training personnel to use AI insights effectively**. Creating models and AI tools that can be easily understood by the end-user is very important – if a Soldier, Airman, Marine, or intelligence analyst cannot easily utilize the tool, it has very little use. Even a highly accurate model can underwhelm if analysts aren't sure how to interpret its outputs or if it isn't integrated into their existing tools. Effective programs have assigned AI "agents" to work alongside analysts, with user-friendly interfaces and clear visualization of why a model flagged something. Over time, human operators grow to trust the AI as they see its utility, and they learn to leverage it to augment their expertise (rather than being replaced by it). This human-machine teaming, where AI handles rote data processing and humans apply judgment, consistently produces better outcomes than either alone. In cyber defense, AI might cluster thousands of alerts into a handful of high-priority incident reports where

human analysts then investigate those summaries. The feedback loop – analysts validating AI findings and feeding results back to refine the model – is key to continuous improvement. Finally, national security agencies understandably demand rigorous testing (including red-teaming and scenario-based exercises) before deploying AI operationally. We've learned the importance of robust testing and validation of AI in these contexts, but each use case and threshold is different, making it essential that personnel understand the capabilities and the explicit workflows they wish to adapt AI to.

AI is also poised to help U.S. cyber and intelligence personnel counter new threats posed by adversarial use of AI. We must assume that competitor nations and malicious actors are already employing AI-enabled cyber capabilities – from automating their intrusion attempts to using AI to generate deceptive deepfakes, more convincing phishing lures and create information warfare. Indeed, the Pentagon's National Defense Strategy warned that state competitors and even non-state actors will have access to emerging technologies like AI, potentially eroding the U.S. military's traditional advantage. In the cyber domain, this means our adversaries **could leverage AI to find software vulnerabilities faster, to tailor attacks that evade detection, or to deploy autonomous malware that adapts on the fly**. To stay ahead of these AI-augmented threats, DoD must likewise incorporate AI across its cyber defenses. For instance, AI can help identify attack toolkits or behaviors that indicate when an adversary might be using machine-driven methods, allowing us to devise countermeasures. AI-enabled threat intelligence can also predict an enemy's next move by analyzing trends, and help our cyber warriors prepare accordingly. AI must be central to the next generation of cyber operations – both offensive and defensive – and it is imperative that DoD embrace these tools to defend our networks, protect critical infrastructure, and maintain decision superiority over any adversary leveraging AI against us.

In summary, AI can greatly enhance pattern recognition, multilingual understanding, and analytic efficiency for defense and intelligence, provided we integrate these tools thoughtfully – with attention to the data they're trained on, the way humans interact with them, and thorough validation to ensure they perform as intended under real-world conditions.

2. Key Technical Considerations for AI in Defense

When evaluating AI solutions for defense applications, several critical factors deserve consideration beyond raw capability.

First, right-sizing models for specific operational contexts should be a priority. The AI industry as of late has been focused on scaling to ever-larger general purpose

models, but defense environments frequently benefit more from specialized, optimized models designed for specific tasks. This means **building the *right* model for the mission**: models that are efficient, adaptable, low latency, and can be deployed in the field under real-world constraints. This has been a major area of focus for Cohere. Smaller, carefully tuned models (in the range of 7 billion parameters, compared to hundreds of billions of parameters) can deliver excellent performance without the enormous computational requirements of much larger systems. This approach allows models to run on **limited hardware** (such as tactical edge devices or classified on-premises servers) and to operate within the latency and power constraints of military use cases. Efforts to improve model efficiency are also key to keeping hardware requirements small for larger, more performant models. Our most recent large language model release, Command A, outperforms the largest state-of-the-art models in many enterprise-specific benchmarks while being 111 billion parameters - just one-sixth the size of similarly performing competitors - and **only requires two GPUs to run** (compared to 32 or more for larger models).

Our adversaries across the world have also been targeting these types of efficiency improvements - as witnessed by the recent release of Deepseek V3. They recognize that the distributive impact of AI will only occur if models are capable of being integrated in workflows regardless of context or compute power. Cohere has always been focused on model efficiency, long before Deepseek was released. Deepseek is impressive, but not groundbreaking: while it was a wake-up call for some, for us it merely validated that our approach has always been the right one.

Second, collaborative model development between government and industry represents another vital priority. The most effective approach allows agencies to leverage commercial AI expertise while maintaining appropriate security control. **DoD does not need to undertake the costly, time-consuming task of developing every AI model from scratch. A partnership model could enable the customization of models to specific operational contexts using domain-specific data while protecting sensitive information.** Such collaboration accelerates deployment while giving agencies confidence in model behavior since systems can be thoroughly tested and validated on relevant data before implementation. Speeding “time-to-value” through these partnerships will allow DoD to bring capabilities to bear faster in solutions that meet military users where they are – in secure facilities, at the tactical edge, and at the speed of mission – rather than expecting DoD to conform to a one-size-fits-all AI built for consumer use.

At Cohere, we work hand-in-hand with our customers to customize or collaborate in building new AI models suited to their specific operational context. Cohere's team is able to take domain-specific data (for example, technical documents, or intel reports)

and train specialized models that capture the nuances of DoD's unique problems. We have done this in the commercial and national security space, working closely to develop custom large language models from proprietary data, and we apply the same approach to the defense market. The result is an AI model that is tailored to DoD's vocabulary, threat landscape, and security constraints – but developed rapidly and cost-effectively through public-private partnership.

Third, secure, flexible deployment architecture is an essential technical consideration for adopting AI in defense. High security deployments typically do not choose public cloud services for sensitive data. Instead, the preference for AI solutions is that data can be hosted in secure environments. Cloud-agnostic and hardware-agnostic approaches ensure AI systems can be deployed in the widest range of these secure environments—from private cloud services to classified data centers to edge computing devices—without creating specific or single vendor dependencies. The DoD needs systems that are **interoperable and work across all cloud and chip types** to prevent vendor lock-in and allow for a hybrid AI ecosystem to meet its needs.

AI systems deployed in defense and national security contexts must also maintain data sovereignty, allowing organizations to retain full control of their information. In the field of artificial intelligence, we believe "**AI sovereignty**" is also necessary in certain cases. This can include several aspects, but we regard it as a solution in which not only the data, but also related AI models are kept within a specific country, on that country's (or customer's) infrastructure. In these "sovereign" or private AI deployments, developers should have no access to the customer's data. This means helping customers confidently customize and deploy models on their own infrastructure.

At Cohere, we are proud to have remained independent. We work across all major cloud systems and even allow private deployments rather than being locked into one cloud provider. This deployment flexibility is crucial for both operational agility and supply chain security. Our models are purpose-built to be hardware agnostic, allowing for flexible deployment of our models across chip technologies, cloud platforms, and even private, air-gapped environments. For private, air-gapped environments, our models are deployed wherever the data is, maintaining the strictest security and privacy standards. This is, in direct contrast to other deployment approaches where sensitive data must be sent to public clouds and wherever AI models are available/served.

Cohere also recognizes that properly securing AI requires going beyond traditional controls. Our Secure AI framework details our holistic approach to managing risk by implementing security safeguards throughout the development lifecycle of an AI model. We believe that protective measures should be designed within the models and also the environment that models are developed in. We take this a step further and work in

partnership with its customers to conduct supplementary assurance evaluations or testing as needed.

3. How DoD’s Existing Technology Constraints Apply to AI Adoption

As the DoD strives to integrate AI across its enterprise, it's worth reflecting on lessons from previous technology modernization efforts and assessing where DoD stands today. A recurring theme is that the biggest obstacles to adoption are often organizational and procedural, rather than purely technical. The DoD has no shortage of pilot projects demonstrating AI's potential, but scaling those successes DoD-wide has proved challenging. Legacy procurement rules, lengthy accreditation processes, siloed data, and workforce gaps can slow down the fielding of new AI capabilities. In many ways, these hurdles echo what we saw in prior IT modernization waves (like cloud computing or mobile tech) – the technology may be ready, but the bureaucracy must catch up. For example, the DoD's Authority to Operate (ATO) process for certifying new software has historically been cumbersome and slow, sometimes taking 12-18 months for a new system. Such delays are ill-suited to AI systems, which often need frequent model updates or iterative deployment. Studies have noted that procedures like the ATO, while important for security, can stifle rapid innovation if not reformed. The DoD has begun addressing these issues by streamlining risk management frameworks and adopting more agile approaches (such as DevSecOps and continuous ATO in some programs), but there is more work to do to truly accelerate AI uptake.

Another lesson is the need to invest in the **enabling infrastructure and talent** that make AI adoption possible. The Department's AI efforts in recent years – from Project Maven to the Joint Artificial Intelligence Center (JAIC) – revealed shortcomings in areas like data management, computing resources, and AI skills among personnel. Put simply, **deploying AI is not just about the algorithm; it requires quality data, reliable data pipelines, modern cloud or edge computing environments, and a workforce trained to develop, test, and use AI.** DoD's Chief Digital and AI Office (CDAO) is actively tackling these foundational issues by promoting data standards, standing up AI development platforms, and expanding training and collaboration opportunities across the Joint Force, all of which are positive steps. The current state of AI integration in DoD features numerous promising prototypes and niche deployments (for instance, AI in predictive maintenance, intel analysis, or business process automation), but enterprise-scale implementation – where AI is a routine part of military operations – is still in progress. Key doctrines and strategies (such as the 2022 DoD Data Strategy and AI Strategy) emphasize being "AI-ready," yet many operational units are still unfamiliar with AI tools. This indicates a need for continued top-down leadership emphasis and resource prioritization to move from experimentation to broad adoption.

Crucially, the DoD must balance innovation with security as it integrates AI. Defense officials often point out that we cannot afford to lag in AI adoption because our adversaries are moving quickly. At the same time, the military cannot "move fast and break things" in the way a Silicon Valley startup might; the stakes are too high. The **balance lies in risk-managed innovation – encouraging experimentation and agile deployment of new AI technologies, while also instituting robust testing, evaluation, and governance**. Senior leaders in the Pentagon acknowledge this cultural shift is needed. The DoD must becme more comfortable with experimental failure as a step toward innovation, noting that traditionally the military prizes perfection and zero-failure, but that mindset can hinder rapid adaptation. In recent years, we have seen progress here: initiatives like AI challenges in multiple Military Service labs, the DIU (Defense Innovation Unit) commercial solutions, and "bake-offs" for AI algorithms are injecting a more innovative culture. Yet, maintaining security is paramount – which means AI systems must be thoroughly vetted for vulnerabilities before deployment, and used in accordance with ethical principles and applicable law. The Department has published the “DoD AI Ethical Principles” and is establishing organizations (like the NSA's AI Security Center) to develop best practices for secure AI adoption. The theme moving forward will be governing AI without strangling it: putting in place guidelines, test frameworks, and oversight so we trust the AI in mission settings, but not imposing such onerous processes that field units give up on trying new AI tools. If DoD and Congress can get this balance right – fostering a culture of innovation with accountability – the United States will maintain both a technological edge and the confidence that our use of AI is safe, ethical, and effective.

4. Policy Recommendations for Congress

Four years ago, Congress’ AI Commission took a comprehensive look at the opportunity of applying AI to national security challenges. Their conclusion then was that by 2025 the foundations for widespread integration of AI across DoD must be in place, and that the Department should allocate at least $8 billion toward AI annually. Public estimates from last year were that DoD is allocating about half the amount recommended by the Commission.

To support the Department of Defense in responsibly accelerating AI adoption, I respectfully offer the following policy recommendations for Congress:

(A) Modernize Procurement Processes to Drive AI Innovation: Perhaps the most important step Congress can take is to modernize DoD's acquisition and procurement models for AI and software. The Secretary of Defense’s March 6 memorandum on software acquisition is a key step that clearly signals “software companies make software and the DoD will buy software from software companies.” Current federal

procurement rules and lengthy contracting cycles often favor large, established defense vendors and can inadvertently create high barriers to entry for innovative AI startups. If the U.S. military wants to have the best technology, then we must ensure the best technology companies can compete to serve them – this includes non-traditional defense suppliers like Cohere. Congress should encourage and empower DoD to use more flexible, agile procurement mechanisms tailored to fast-evolving tech. For example, expanding the use of challenge-based solicitations, pilot programs, Commercial Solutions Openings (CSOs) and Other Transaction Authorities (OTAs) can allow the Department to rapidly test and integrate cutting-edge AI solutions. These approaches lower the burden to get in the door, enabling smaller firms with innovative ideas to prove their value through prototypes or competitions rather than needing an extensive track record of past contracts. Alongside this, Congress can push for simplified compliance requirements and streamlined vendor qualification criteria for AI software, so that well-intentioned rules do not unintentionally exclude agile startups.

By making federal procurement more accessible and agile, the government can leverage its "power of the purse" to catalyze AI innovation and avoid stagnation. Importantly, modernizing procurement isn't just about speed – it's also about outcomes. I recommend that Congress direct DoD to **update its source selection criteria for AI-related contracts to place appropriate weight on innovation, security, and performance,** rather than defaulting to corporate size or longest past performance. The goal should be to reward true technical merit and risk mitigation, not just familiarity. Taken together, these procurement reforms will help ensure that the U.S. military has access to the full marketplace of AI innovation – bringing the best capabilities to our cyber warfighters, from large defense primes to innovative startups.

(B) Promote Interoperability and Avoid Vendor Lock-In: As the DoD acquires AI tools from various sources, Congress should ensure that interoperability and open standards are a priority. No single vendor should be able to supply all of DoD's AI needs. The AI ecosystem is dynamic and evolving, and so should be AI vendors selling to the government. We want an ecosystem where multiple AI solutions can plug into defense systems seamlessly and even augment each other. To that end, Congress could require that DoD include interoperability requirements in AI procurements and programs, **mandating that solutions adhere to common data formats, APIs, or integration standards.** This will allow systems from different providers – say an AI cybersecurity sensor from one company and an AI analytics dashboard from another – to work together without heroic integration efforts. It also prevents proprietary lock-in, where DoD might be stuck with one vendor's ecosystem.

Open, interoperable approaches spur competition and innovation, because vendors know their products must operate in a hybrid environment and can be replaced with

improved technology. Alongside interoperability, Congress should encourage DoD to avoid single-source dependency in critical AI capabilities. This might involve funding diverse pilots for a given use-case (so multiple solutions are evaluated), and then adopting the best – or even multiple – solutions. Legislative report language could underscore that avoiding vendor lock-in is a strategic imperative for both security and negotiating power. In summary, Congress should help DoD obtain AI systems that are as flexible and interchangeable as possible. This not only ensures our forces get the best tech, but it also safeguards against supply-chain risks and empowers DoD to rapidly upgrade components as the technology advances.

(C) Support AI Adoption with Internal Benchmarking: While speeding up AI adoption, Congress must also ensure that appropriate benchmarking guidelines are in place for AI in national security. In the commercial space, we have quickly learned that general academic benchmarks are regularly gamed and often don't represent real world use. For example, they don't showcase how models will tackle tasks on an assembly line or in analyzing research. We recommend that Congress and the DoD consider funding programs that develop methods to test and validate AI systems for the kinds of reliability and use cases they will require - including human evaluations.

Conclusion

AI has the potential to significantly strengthen U.S. national defense – but realizing that potential requires deliberate action and partnership. Cohere is fully committed to working with the DoD, Congress, and our allies to advance AI capabilities in cybersecurity and defense. We bring not only technology, but a shared sense of mission. We understand that the stakes are high: the security of the nation and the safety of its service members depend on us getting this right. That is why we will continue prioritizing AI development with rigorous attention to security, privacy, and operational effectiveness. We will continue to innovate on ways to make AI more efficient, more adaptable, and easier to deploy securely at scale.

I urge Congress to take a forward-thinking approach in crafting AI policies for national security – one that balances our national security imperatives with America's spirit of innovation. By modernizing procurement, championing interoperability, insisting on benchmarking based on use case, and tapping into the best insights from research, Congress can accelerate the DoD's adoption of AI in a manner that is both rapid and responsible. In doing so, we can maintain our competitive edge in an era where AI capabilities will increasingly define the strength of our defense. Cohere stands ready to assist in this effort. Working together – industry, DoD, and Congress – I am confident we can build AI systems that bolster our security, reflect our values, and earn the trust of those who rely on them.

Thank you for the opportunity to testify today. I look forward to your questions and to continuing the dialogue on how we can ensure the United States and its allies lead in the secure and responsible use of AI for our national defense.