ANALYSIS

The Good, Bad, and Really Weird AI Provisions in the Annual Defense Policy Bill

Congress is taking steps to mitigate risks to national security and basic rights, but the industry's role needs more scrutiny.

Amos Toh

December 15, 2025

Government Power

PART OF

Artificial Intelligence and National Security [link-1]

MORE ON

Privacy Rights [link-2]

Artificial Intelligence [link-3]

Congress is on the verge of passing the National Defense Authorization Act (NDAA), the annual bill that sets the budget for the Department of Defense (DoD) and other defense and intelligence-related activities across the

government. The NDAA shapes the overall direction of national security investments, and this year's law will have an impact on how artificial intelligence is used in warfighting and intelligence activities for decades to come. The more than 3,000 page law covers a lot of ground, from establishing a steering committee of senior national security officials to set policy on artificial general intelligence [link-4], to prohibiting intelligence agencies from using the Chinese AI model, DeepSeek [link-5]. Here's a breakdown of how Congress is grappling with some of AI's more consequential risks to national security and fundamental rights — what it does well, what's just plain weird, and what it overlooks at the nation's peril.

The good

AI governance

It's a good thing that Congress is finally moving to regulate how intelligence agencies acquire and deploy AI. Agencies have been expanding their reliance on the technology to investigate drug trafficking networks [link-6], surface potential targets [link-7] for military strikes, and scour social media [link-8] for threats of violence or terrorism. These applications pose grave risks to privacy and civil liberties, enabling agencies to collect, piece together and analyze more personal and sensitive data than ever before and generate insights into people's relationships, movements, habits, and political views. The technology can also be inaccurate and misleading [link-9]: shortcomings in how models parse linguistic context and nuance, particularly in non-English languages [link-10], could lead to the misidentification of certain individuals or groups as security threats while overlooking genuine indicators of illicit activity.

This year's NDAA establishes a high-level overview of how agencies should manage these risks. Section 6602 requires heads of intelligence agencies to "track and evaluate performance of procured and [agency]-developed artificial intelligence, including efficacy, safety, fairness, transparency, accountability, appropriateness, lawfulness, and trustworthiness." Where agencies use AI models that are widely

available to the public, such as OpenAI's ChatGPT or Google's Gemini, sec. 6603 requires the development of testing standards and benchmarks along similar criteria.

This framework is threadbare compared to guidance the Biden White House issued in 2024 [link-11], which stipulated the risks national security agencies should mitigate (such as "unlawful discrimination" and "possible failure modes") and the means by which to do so (cost-benefit assessments, evaluations by an "independent reviewing authority," training, whistleblower mechanisms). But the fate of the Biden era guidance — known as National Security Memorandum 25 [link-12] — is uncertain. NSM-25 also granted agencies sweeping discretion [link-13] to waive compliance with its safeguards — loopholes that Congress, for now, neither closes nor endorses. What this year's NDAA affirms, however, is that governance of national security AI is not solely a matter of executive prerogative, but also a legislative requirement. Congress should build on this to develop more detailed safeguards and exercise independent scrutiny and oversight.

Will these safeguards meaningfully reduce error, abuse, and harm, or will they just become another box-checking exercise? This depends on whether general and subjective criteria (such as "efficacy" and "safety") are translated into concrete specifications, how these specifications are integrated into agency workflows, and whether agencies are incentivized to value meaningful accountability rather than rely on superficial compliance. Here, there is a lot more work to be done.

Mitigating outsourcing risks

Lawmakers are also signaling concern about the government's growing dependence on the tech industry for critical national security functions. The federal government already relies substantially on industry for their tech and AI needs, and the intelligence community is no [link-14] exception [link-15]. Recognizing this reality, sec. 6602 directs the intelligence community to develop model contract terms that "address technical data rights,"

"minimize dependency on proprietary information" and prioritize "interoperability" (i.e. software that is easier to link up with existing systems or systems by different vendors). But it stops short of requiring these terms, raising doubts about how strongly agencies will insist on them in contract negotiations.

Nonetheless, strengthening the government's hand to negotiate for more robust data and software rights is a hedge against excessive dependence on tech companies. It gives agencies more wiggle room to break a contract should a vendor under-perform or pose safety concerns. This is critical, since a handful of players already dominate key markets for intelligence and defense AI.

Palantir leads the market [link-16] for helping agencies organize and make sense of their data. Companies such as OpenAI and Google are the only US firms with the data, resources, and talent [link-17] to build foundation models — a form of AI trained on large volumes of data to perform a wide range of tasks, from answering questions to generating video. These forms of market dominance raise a host of national security risks that consolidation within the legacy defense industry has foreshadowed. A 2022 Pentagon study [link-18], for example, found that the lack of competition can lead to shortages of critical equipment and hamper innovation and performance in government contracts.

Another notable restriction is sec. 1654, which prohibits the Secretary of Defense from outsourcing "kinetic missile defense capabilities" to a "subscription-based service, a pay-for service model, or a recurring-fee model to engage or intercept a target." It also emphasizes that the decision to engage these capabilities — "including targeting, launch authorization, and engagement of airborne or spaceborne threats" — is an "inherently governmental function" that only federal employees or the armed forces can perform.

This section appears to rebuff attempts by Space X and other defense tech companies vying for the "Golden Dome" project [link-19] to pitch the Trump administration [link-20] on a "subscription" model to missile defense, where the companies will own the

system and license its use to the Pentagon. This has sparked broader fears that the military will lose control of the development of critical defense systems and associated costs, while empowering contractors to shut off or restrict these systems for arbitrary reasons [link-21] .

But the NDAA fails to address the spread of the subscription-based model throughout a wide range of surveillance and targeting functions. The Maven Smart System — an AI-based mission control system the Army uses to analyze satellite imagery, drone footage, and other data collected from military sensors — is owned by Palantir and licensed to the Army [link-22] . While this form of outsourcing may be convenient and cost effective in the short term, it disincentivizes the military [link-23] from developing the technical infrastructure and know-how to conduct key aspects of intelligence analysis – and understand and mitigate its risks.

One category of risk lies in cybersecurity. AI and safety engineering experts have warned that the integration of commercially developed AI into military systems increases the range of security vulnerabilities [link-24] adversaries can exploit to undermine these systems. Section 1513 of the NDAA requires DoD to establish a framework for mitigating key risks, such as the possibility that data used to train AI models could be compromised, tampered with, or stolen. Unfortunately, it is silent on a looming danger to privacy: that the personal data that tech companies train their models on can be reconstructed [link-25] by intelligence and defense agencies downstream, as well as by adversaries that gain access to these models.

Oversight of autonomous weapons

The NDAA takes a small but critical step towards improving oversight of autonomous and semi-autonomous weapons. Section 1061 of the NDAA now requires the Pentagon to report waivers of DoD Directive 3000.09 — its internal safeguards for developing such weapons — to the congressional defense committees. This must include descriptions of weapons systems covered by waivers, as well as their rationale and anticipated duration.

The Directive's most stringent safeguard — review and sign off by senior Pentagon leaders before development — is reserved for autonomous weapons, which is defined as systems that can "select and engage targets without further intervention by an operator." These may, for example, include drones programmed to identify and fire on military installations without requiring operator confirmation between identification and firing. While the NDAA does not limit the scope of these waivers, it at least provides a measure of transparency.

The weird

The NDAA's provisions on AI governance for the intelligence community establish a strange "rule of construction": intelligence agencies should not interpret these standards as authority to direct vendors to "alter a model to favor a particular viewpoint."

This seems to be inspired by the Trump administration's effort to weed out "woke AI" in federal procurement. In an executive order issued [link-26] in July 2025, the administration directed agencies to withhold contracts from companies that don't align their technology with "truth-seeking" and "ideological neutrality." I and others have [link-27] argued [link-28] that compliance with the order will likely require companies to degrade AI performance and undermine public trust in their technology. The administration's hostility to diversity, equity, and inclusion, for example, could lead companies to pull back on bias testing and other protocols to measure and account for their technology's impacts on minorities and other underserved communities. But these measures are widely regarded as essential to improving the accuracy of AI systems.

Perhaps recognizing that this thinly veiled attempt to bias AI in favor of the administration's versions of truth and reality may inadvertently compromise the technology, the order permits exceptions for national security "as appropriate." The NDAA changes little. It is doubtful that intelligence agencies would violate its rule of construction if they direct vendors to conduct bias testing and mitigation and turn over the results — these measures may after all ensure that AI models generate more

accurate and comprehensive summaries of contested views. In any event, there is nothing to prevent agency leaders from invoking their existing authority [link-29] to test and evaluate AI performance to require these measures.

The really bad

Where the NDAA takes a disastrous turn is its gutting of safeguards designed to protect taxpayer dollars in military contracting. Some lawmakers have framed these changes [link-30] as a long overdue effort to cut red tape and promote innovation. But the delays and cost overruns that have become endemic in defense acquisition can often be traced [link-31] to the *lack* of sufficient pricing controls and testing early in the process — a problem the NDAA will make worse.

Pentagon fraud and waste is persistent and staggering — some estimate that this has cost taxpayers hundreds of billions of dollars [link-32]. In the 1960s, Congress tried to fix some of the waste by requiring contractors [link-33] to provide data about why their technology costs as much as they say it does, and certify this data as "accurate, complete, and current [link-34]." This is especially important in contracts where the Pentagon has few ways to compare costs and pricing — such as when it relies on a single supplier for key military technology, equipment or parts. These requirements also come with [link-35] remedies for defective pricing — namely, a price adjustment [link-36] if the Pentagon discovers that a contractor legally mandated to certify its data knowingly submitted inaccurate, incomplete, or old cost and pricing information.

The military is reliant on just a handful of contractors — and sometimes only just one [link-37] — for both its AI capabilities and the infrastructure needed [link-38] to support their rollout. And while tech companies are pitching AI as a cheaper alternative to traditional weapons, costs can easily spiral out of control [link-39]. Pricing transparency empowers the Pentagon to counter price gouging and keep costs in check before it's too late.

The NDAA escalates a decades-long effort [link-40] to roll back data disclosures that keep a lid on

overcharging. The General Accountability Office estimates [link-41] that less than one percent of DoD and NASA contracts require certified cost and pricing data. The NDAA will exempt even more contracts from this requirement. Previously, only qualifying contracts worth $2 million or more were subject to mandatory disclosure; sec. 1804 raises this to $10 million. Sec. 1826 would exempt "nontraditional defense contractors" from certified data requirements entirely — a category that would include many firms in the tech industry.

This push is part of a broader effort to fast-track public-private partnerships to modernize critical defense systems. But lawmakers are overlooking threats to civil liberties and national security when the government cedes too much control of these systems to Silicon Valley. The cybersecurity risks Congress is concerned about, for example, are showing up in a battlefield communications system [link-42] that Anduril, Palantir and other tech companies are building for the Army. A memo from the Army's chief technology officer found that the system is a "black box," and that the service can't control who has access to what. The Army now says [link-43] that these deficiencies have been mitigated.

This reflects the NDAA's contradictory impulses on AI. On one hand, Congress is trying to take much-needed steps to mitigate the technology's risks to national security and fundamental rights. But it is also pushing defense and intelligence agencies towards faster and greater dependence on industry, increasing exposure to these risks.

Overall, the NDAA represents a significant codification of US AI policy in an important domain. It has made welcome progress, however slight, towards addressing risks with the technology. But its failure to right-size industry's role in developing this technology for national security purposes could have lasting repercussions.

This article first appeared at Tech Policy Press. [link-44]

Links

link-1: Artificial Intelligence and National Security /series/artificial-intelligence-and-national-security

link-2: Privacy Rights https://www.brennancenter.org/topics/government-power/privacy-rights

link-3: Artificial Intelligence /tags/artificial-intelligence

link-4: artificial general intelligence https://www.techpolicy.press/artificial-general-intelligence-what-are-we-investing-in/

link-5: DeepSeek https://www.nist.gov/news-events/news/2025/09/caisi-evaluation-deepseek-ai-models-finds-shortcomings-and-risks

link-6: drug trafficking networks https://www.pbs.org/newshour/world/u-s-intelligence-agencies-embrace-of-generative-ai-is-at-once-wary-and-urgent

link-7: surface potential targets https://www.bloomberg.com/features/2024-ai-warfare-project-maven/

link-8: scour social media https://www.brennancenter.org/our-work/research-reports/social-media-surveillance-us-government

link-9: inaccurate and misleading https://www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html

link-10: non-English languages https://cdt.org/insights/lost-in-translation-large-language-models-in-non-english-content-analysis/

link-11: issued in 2024 https://perma.cc/S747-YLWA

link-12: known as National Security Memorandum 25 https://perma.cc/S747-YLWA

link-13: sweeping discretion https://www.lawfaremedia.org/article/narrowing-the-national-security-exception-to-federal-ai-guardrails

link-14: no https://www.nextgov.com/artificial-intelligence/2024/09/cia-looks-fast-track-ai-adoption-through-cloud-contract/399278/

link-15: exception https://breakingdefense.com/2025/11/startup-enabled-intelligence-nabs-ngas-708-million-ai-training-contract/

link-16: leads the market https://www.nytimes.com/2025/11/14/business/palantir-ai-military-stock-market-.html

link-17: data, resources, and talent https://ainowinstitute.org/publications/compute-and-ai

link-18: A 2022 Pentagon study https://media.defense.gov/2022/feb/15/2002939087/-1/-1/1/state-of-competition-within-the-defense-industrial-base.pdf?inline=1

link-19: "Golden Dome" project https://www.congress.gov/crs-product/IF13115

link-20: the Trump administration https://www.reuters.com/business/aerospace-defense/musks-spacex-is-frontrunner-build-trumps-golden-dome-missile-shield-2025-04-17/

link-21: arbitrary reasons https://www.reuters.com/investigations/musk-ordered-shutdown-starlink-satellite-service-ukraine-retook-territory-russia-2025-07-25/

link-22: licensed to the Army https://perma.cc/8H4W-DQ6N

link-23: disincentivizes the military https://www.nybooks.com/online/2025/10/04/the-war-over-defense-tech/

link-24: increases the range of security vulnerabilities https://ainowinstitute.org/wp-content/uploads/2024/10/Mind_the_Gap__Foundation_Models_and_the_Covert_Proliferation_of_Military_Intelligence__Surveillance__and_Targeting.pdf?inline=1

link-25: can be reconstructed https://arxiv.org/abs/2405.20272

link-26: issued https://www.whitehouse.gov/presidential-actions/2025/07/preventing-woke-ai-in-the-federal-government/

link-27: have https://www.brennancenter.org/our-work/analysis-opinion/how-trumps-ai-policy-could-compromise-technology

link-28: argued https://www.techpolicy.press/trumps-order-against-woke-ai-will-create-real-harm/

link-29: their existing authority https://www.law.cornell.edu/uscode/text/50/3334m

link-30: framed these changes https://armedservices.house.gov/uploadedfiles/fy26_ndaa_conference_text_legislative_summary.pdf?inline=1

link-31: often be traced https://www.lawfaremedia.org/article/how-acquisition-reform-could-make-military-ai-more-expensive-and-less-safe

link-32: hundreds of billions of dollars https://www.theguardian.com/commentisfree/2025/mar/15/doge-pentagon-defense-spending

link-33: requiring contractors https://natlawreview.com/article/truth-negotiations-act

link-34: accurate, complete, and current https://www.law.cornell.edu/cfr/text/48/2.101#:~:text=Certified%20cost%20or%20pricing%20data,chapter%2035).

link-35: requirements also come with https://www.gao.gov/assets/gao-22-105307.pdf?inline=1#page=11

link-36: a price adjustment https://www.acquisition.gov/far/15.407-1

link-37: sometimes only just one https://breakingdefense.com/2024/05/new-contract-expands-maven-ais-users-from-hundreds-to-thousands-worldwide-palantir-says/

link-38: infrastructure needed https://www.warren.senate.gov/newsroom/press-releases/warren-schmitt-renew-bipartisan-fight-for-more-competition-in-pentagons-ai-and-cloud-contracting

link-39: spiral out of control https://www.forbes.com/sites/katharinabuchholz/2024/08/23/the-extreme-cost-of-training-ai-models/

link-40: decades-long effort https://www.stimson.org/2025/acquisition-reform-at-a-crossroads/

link-41: estimates https://www.gao.gov/products/gao-22-105307

link-42: battlefield communications system https://www.reuters.com/business/aerospace-defense/anduril-palantir-battlefield-communication-system-has-deep-flaws-army-memo-says-2025-10-03/

link-43: now says https://breakingdefense.com/2025/10/army-says-its-mitigated-critical-cybersecurity-deficiencies-in-early-ngc2-prototype/

link-44: Tech Policy Press. https://www.techpolicy.press/the-good-bad-and-really-weird-ai-provisions-in-the-annual-us-defense-policy-bill/