



ARTIFICIAL INTELLIGENCE IN NUCLEAR COMMAND, CONTROL & COMMUNICATIONS: A TECHNICAL PRIMER

SYLVIA MISHRA

PHILIP REINER

SEPTEMBER 2025



**Artificial Intelligence in Nuclear Command, Control & Communications:
A Technical Primer**

September 2025

Authors: Sylvia Mishra and Philip Reiner

Design: Taylor White

The Institute for Security and Technology and the authors of this report invite free use of the information within for educational purposes, requiring only that the reproduced material clearly cite the full source.

Copyright 2025, The Institute for Security and Technology
Printed in the United States of America



About the Institute for Security and Technology

Uniting technology and policy leaders to create actionable solutions to emerging security challenges

Technology has the potential to unlock greater knowledge, enhance our collective capabilities, and create new opportunities for growth and innovation. However, insecure, negligent, or exploitative technological advancements can threaten global security and stability. Anticipating these issues and guiding the development of trustworthy technology is essential to preserve what we all value.

The Institute for Security and Technology (IST), the 501(c)(3) critical action think tank, stands at the forefront of this imperative, uniting policymakers, technology experts, and industry leaders to identify and translate discourse into impact. We take collaborative action to advance national security and global stability through technology built on trust, guiding businesses and governments with hands-on expertise, in-depth analysis, and a global network.

We work across three analytical pillars: the **Future of Digital Security**, examining the systemic security risks of societal dependence on digital technologies; **Geopolitics of Technology**, anticipating the positive and negative security effects of emerging, disruptive technologies on the international balance of power, within states, and between governments and industries; and **Innovation and Catastrophic Risk**, providing deep technical and analytical expertise on technology-derived existential threats to society.

Learn more: <https://securityandtechnology.org/>

Acknowledgments

This report emerged from numerous insightful discussions and exchanges of ideas with experts who emphasized the need for better understanding of AI terminologies, what they stand for, and their implications for nuclear systems.

The authors extend their sincere gratitude to Malo Bourgon, CEO of the Machine Intelligence Research Institute (MIRI), for his invaluable feedback, comprehensive review, and recommendations that significantly strengthened this report. We would also like to thank our colleagues Kathryn Harris and Sandip Sarkar at Scale AI for their thoughtful commentary and substantive engagement with our research exploring the critical intersection of artificial intelligence and national security. Their expertise and insights have been instrumental in refining our analysis and conclusions.

We express our deepest gratitude to Peter Hayes, Herbert Lin, and Paul Scharre for their expert insights on the implications of AI on nuclear weapons systems.

This work would not have been possible without the unwavering support and encouragement of Matthew Gentzel and Carl Robichaud at Longview Philanthropy. We are deeply grateful for their generous funding and continued commitment to advancing research at the vital nexus of AI and Nuclear Command, Control, and Communications (NC3). Their vision and support have enabled us to tackle these complex challenges at the forefront of emerging technologies and nuclear policy.

The views and conclusions expressed in this report remain solely those of the authors and do not necessarily reflect the positions of the individuals or organizations mentioned above.

Contents

- Introduction1**
- Current U.S. Strategic Position 3**
- AI and its Historical Role in NC3..... 4**
- Other broad characteristics of Narrow AI tools applied in weapons systems:..... 5
- Historical AI Use Cases in NC3 Systems..... 6**
- Potential Limitations and Challenges..... 8**
 - Enterprise 9*
 - Military.....9*
- Glossary of Terms10**

Introduction

Automation and artificial intelligence (AI) have been integral to nuclear weapons systems and operations for decades.¹ Given AI-enabled tools' ability to assess enormous amounts of intelligence data at unprecedented speeds, the integration of modern AI machine learning programs with nuclear command, control, and communications (NC3) systems will only grow with time.² On April 10, 2025, the Institute for Security and Technology (IST) held a workshop entitled, "The Risks and Opportunities of AI in NC3: Finding Common Ground" in Washington, D.C. The workshop explored how the integration of novel AI into global NC3 systems over the next five years could transform strategic stability and deterrence dynamics.

To better understand what constitutes 'novel' AI in the context of nuclear weapons decision-making, this primer examines which AI tools are currently being used in NC3 systems, and how we should consider the likely technological opportunities that are forthcoming with advanced AI capabilities. This primer is not exhaustive, but aims to capture and indicate ways in which AI tools and technologies have been integrated and adapted into NC3 systems and what a potential future will portend.

For the purposes of this workshop, to be specific and solely focused on narrow/specialized AI, the authors of the report deliberately did not consider the impact of artificial general intelligence (AGI) or artificial superintelligence (ASI). While the potential for those technical advances is real, as are their implications for NC3, these were not the focus of the April workshop.

Workshop participants were presented with this primer, and encouraged to consider AI as defined along the following parameters:

The state of the art in AI has evolved from narrow task-specific systems to more flexible, generalized, and agentic tools that can plan, reason, and execute complex workflows autonomously. In enterprise settings, multi-agent systems are being deployed to augment executive decision-making by analyzing vast datasets, generating scenarios, and providing evidence-based recommendations. These systems are increasingly capable of handling complex reasoning chains and maintaining context across extended operations. In military applications, both Israel and Ukraine have deployed AI-powered systems for intelligence analysis, autonomous reconnaissance, and command-and-control optimization. The Israeli Defense Forces' "Gospel"³ system processes battlefield data for tactical recommendations, while Ukraine has leveraged AI

1 Vladislav Chernavskikh, "Nuclear Weapons and Artificial Intelligence: Technological Promises and Practical Realities," SIPRI Background Paper, September 2024, <https://doi.org/10.55163/VBQX6088>.

2 James Johnson, *AI and the Bomb: Nuclear Strategy and Risk in the Digital Age*, (Oxford University Press, May 2023).

3 Anna Aronheim, "Israel's operation against Hamas was the world's first AI war," *Jerusalem Post*, May 27, 2021, <https://www.jpost.com/arab-israeli-conflict/gaza-news/guardian-of-the-walls-the-first-ai-war-669371>.

for drone operations and countering information operations.⁴ Improvements in the capabilities and utility of more general AI systems have been marked by better multimodal understanding, longer context windows, and better tool integration. Within the next 4 to 5 years, these systems will likely advance in reasoning capabilities, multimodal integration, and adversarial robustness, potentially enabling more autonomous planning and faster operational tempo in military contexts, while maintaining human oversight. As these systems become more capable and more autonomous, important questions arise: **what constitutes responsible oversight, and how responsible is that oversight?**

Military leaders can expect AI systems to evolve into highly integrated decision support platforms providing unprecedented situational awareness, predictive intelligence, and operational planning capabilities. Advanced multi-agent systems will likely enable real-time integration of vast intelligence sources, automatically identifying patterns that human analysts might miss, while generating and evaluating multiple courses of action in seconds. Autonomous systems will operate collaboratively in contested environments with minimal supervision, rapidly adapting to changing conditions. Command interfaces will become more intuitive, with natural language processing and mixed-reality visualizations making complex battlefield information immediately accessible to commanders. These developments will potentially enable smaller forces to achieve greater effects through superior information advantage, compressed decision cycles, and precision resource allocation—all while maintaining appropriate human judgment over critical decisions.

Multi-agent systems will likely approach exceptional sophistication. These systems will feature specialized agents with domain-specific expertise collaborating through structured protocols and shared representational frameworks. Their capabilities will include complex multi-step reasoning across extended time horizons, robust operation in information-limited environments, and seamless integration of multimodal inputs spanning visual, textual, and sensor data. Architectures will emphasize both delegation mechanisms for task decomposition and coordination frameworks that balance agent autonomy with system-wide coherence.

Scaling laws will continue to drive capability improvements, with larger parameter counts enabling more sophisticated reasoning and representation. However, architectural innovations are increasingly complementing raw scale, with mixture-of-experts approaches, sparse activation patterns, and specialized routing mechanisms allowing more efficient computation. Test-time scaling through chain-of-thought reasoning and reinforcement learning on reasoning processes has emerged as another complementary approach, where models apply deliberate step-by-step thinking to improve inference without increasing parameter counts. Retrieval-augmented generation techniques will expand context utilization beyond traditional context window constraints, while agent-specific memory systems will enable persistent understanding across extended operations (see [Glossary of Terms](#) for further clarification on terminology).

4 Sam Bendett and David Kiricheko, “Battlefield Drones and the Accelerating Autonomous Arms Race in Ukraine,” Modern War Institute at West Point, January 1, 2025, <https://mwi.westpoint.edu/battlefield-drones-and-the-accelerating-autonomous-arms-race-in-ukraine/>.

Current U.S. Strategic Position

The recently released United States Strategic Command 2025 Posture statement clearly articulates that the United States will maintain a human “in the loop”:

“STRATCOM will use Artificial Intelligence/Machine Learning (AI/ML) to enable and accelerate human decision-making. To fully utilize the potential of AI, USSTRATCOM requires data scientists with expertise in AI and advanced platforms across multiple classifications. Opportunities exist to leverage the emerging digital engineering environment to bridge the gap toward adopting AI/ML into the nuclear systems architecture. AI will remain subordinate to the authority and accountability vested in humans. In all cases, the United States will maintain a human ‘in the loop’ for all actions critical to informing and executing decisions by the President to initiate and terminate nuclear weapon employment.”⁵

General Anthony Cotton, Commander of U.S. Strategic Command, has emphasized that while AI and data analytics can enhance decision-making and improve deterrence, human control over nuclear weapons decisions remains paramount, and AI should never make those decisions autonomously.⁶ Together, these statements indicate steady and potentially accelerating integration of AI technologies into U.S. military operations. We have already seen an evolution in how STRATCOM and other Department of Defense elements discuss AI. Today, we see greater openness and willingness to admit the need for integrating cutting-edge AI into NC3. This is a significant change from 2018, when the annual Nuclear Posture Review contained no direct mention of AI, merely outlining, “the United States will continue to adapt new technologies for information display and data analysis to improve support for Presidential decision making and senior leadership consultations.”⁷ This progression from a complete absence of AI references to STRATCOM’s explicit requirements for AI to ‘enable’ and ‘accelerate decision making’ marks a crucial shift that demonstrates how AI integration with NC3 systems is not only taking shape, but also gaining momentum.

5 *Hearing on United States Strategic Command and United States Space Command in review of the Defense Authorization Request for Fiscal Year 2026, Before the Subcommittee on Strategic Forces of the Committee on the Armed Services, 119th Cong. 1, 2025, statement of General Anthony J. Cotton, Commander of U.S. Strategic Command, https://www.armed-services.senate.gov/imo/media/doc/testimony_of_general_anthony_jcotton2.pdf.*

6 *Hearing on United States Strategic Command and United States Space Command in review of the Defense Authorization Request for Fiscal Year 2026.*

7 Department of Defense, “Nuclear Posture Review,” February 2018, <https://media.defense.gov/2018/feb/02/2001872877/-1/-1/1/executive-summary.pdf>

AI and its Historical Role in NC3

It is well understood that the general advancement and pace of AI technology continues to accelerate. Currently, most nuclear-armed states maintain policies requiring human judgment and authorization in nuclear weapons policy making.⁸ The AI systems being integrated into NC3 are narrow in scope, designed for specific functions rather than general-purpose reasoning. This type of AI can learn from repetitive tasks and their environment in order to improve their performance.⁹ These modular or narrowly defined AI tools are being utilized for support roles and enabling roles, such as monitoring, detection, analysis and translation, predictive maintenance, improving safety margins, streamlining operations, real-time monitoring, and operator training, among other applications.¹⁰ Such AI tools operate within tightly defined parameters and rules of engagement. When discussing AI integration in NC3 systems to date, we are thus still primarily referring to narrow or specialized AI.

Narrow or Specialized AI: Narrow or specialized AI are systems designed to perform specific tasks within well-defined parameters. These systems excel at particular functions like pattern recognition, anomaly detection, or predictive analysis within their domains of expertise. Due to their ability to process large amounts of data and identify patterns, narrow AI systems excel at specific tasks. Narrow AI models aren't able to transfer knowledge from one domain to another, or understand the broader context of their actions.¹¹ Narrow AI systems are specialized for bounded tasks and may use supervised, unsupervised/self-supervised, reinforcement learning, or rule-based approaches, often with large datasets when available, but not inherently reliant on them. Narrow AI also requires extensive training and often needs retraining for new tasks or changes in its environment.¹² Some examples of narrow AI are Expert Systems, Decision Trees, and Natural Language Processing. Expert systems equipped with Narrow AI capabilities can be trained to emulate the human decision-making process and apply expertise to solve complex problems.¹³

8 Jill Hurby and M. Nina Miller, "Assessing and Managing the Benefits and Risks of Artificial Intelligence in Nuclear-Weapon Systems," Nuclear Threat Initiative, August 2021, <https://www.nti.org/analysis/articles/assessing-and-managing-the-benefits-and-risks-of-artificial-intelligence-in-nuclear-weapon-systems/>.

9 Alexandra Tsitsirigos, "Narrow Artificial Intelligence Weapons Systems and their Impact on the Balance of Power," *Yale Review of International Studies* (March 2020, Winter Issue), https://yris.yira.org/essays/narrow-artificial-intelligence-weapons-systems-and-their-impact-on-the-balance-of-power/#_ftn5.

10 Alexa Wehsener, Andrew W. Reddie, Leah Walker, and Philip Reiner, "AI-NC3 Integration in an Adversarial Context: Strategic Stability Risks," Institute for Security and Technology, February 2023, <https://securityandtechnology.org/virtual-library/reports/ai-nc3-integration-in-an-adversarial-context-strategic-stability-risks-and-confidence-building-measures/>.

11 SingularityNET, "A Deep Dive on the Differences between Narrow AI and AGI," Medium, June 15, 2024, <https://medium.com/singularitynet/a-deep-dive-on-the-differences-between-narrow-ai-and-agi-19016011c966>.

12 SingularityNET, "A Deep Dive on the Differences between Narrow AI and AGI."

13 "Understanding the Different Types of Artificial Intelligence," IBM, October 12, 2023, <https://www.ibm.com/think/topics/artificial-intelligence-types>.

Other broad characteristics of Narrow AI tools applied in weapons systems:

Traditional Machine Learning: Traditional machine learning is a method where algorithms learn from a given set of data, drawing patterns and making decisions based on predefined features. It requires significant expertise in selecting and designing these features to optimize performance. Many systems use conventional statistical approaches, supervised learning, and rules-based algorithms rather than the more recent large language model architectures. Missile early-warning systems rely on rule-based AI to identify the launch and trajectory of ballistic missiles with space- or ground-based sensors and transmit this information to human operators for validation.¹⁴ The subset of machine learning we are considering today, deep learning (DL), uses multilayered neural networks, also known as deep neural networks, to simulate the complex decision-making power of the human brain. While traditional machine learning models use simple neural networks with one or two computational layers, DL models use hundreds or thousands of layers to train models.¹⁵ Rapid advances in AI capabilities, especially DL models, have enhanced the performance of AI systems for tasks that require recognizing complex patterns, such as computer vision for classifying objects, people or scenes, natural language processing, and signal recognition (e.g. acoustic or electromagnetic signatures).¹⁶

Deterministic Systems: A deterministic AI system operates under the principle that every set of particular conditions will always lead to the same outcome—one in which the same input will always produce the same output.¹⁷ Examples of deterministic AI systems include Rules-Based Systems, Pathfinding Algorithms, and Expert Systems. When it comes to AI integration with nuclear weapons systems, security planners have historically prioritized systems whose reasoning processes are transparent, predictable, and verifiable.¹⁸

14 Chernavskikh, “Nuclear Weapons and Artificial Intelligence.”

15 Jim Holdsworth and Mark Scapicchio, “What is Deep Learning?” IBM, June 17, 2024, <https://www.ibm.com/think/topics/deep-learning>.

16 Chernavskikh, “Nuclear Weapons and Artificial Intelligence.”

17 “Probabilistic and Deterministic Results in AI Systems,” Gaine Technology, June 31, 2023, <https://www.gaine.com/blog/probabilistic-and-deterministic-results-in-ai-system>

18 “Probabilistic and Deterministic Results in AI Systems,” Gaine Technology.

Historical AI Use Cases in NC3 Systems

Machine Learning Algorithms: These algorithms analyze historical data to identify patterns that predict equipment and system failures, enabling proactive maintenance and minimizing critical downtime. With continuous real-time monitoring capabilities, AI systems can identify anomalies, alert human operators to potential issues, enhance situational awareness, and enable more rapid responses to emerging threats. Researchers at the U.S. Army's Combat Capabilities Development Command's Army Research Laboratory trained a number of classical machine learning algorithms to operate in constrained environments, particularly those involving coalitions.¹⁹ Machine learning systems produce significant benefits by recognizing potential attack signatures and ensuring accurate, timely information reaches decision-makers.

For instance, these systems enhance the already increasingly sophisticated data being relayed to the North American Aerospace Defense (NORAD) Command in Colorado as part of the early-warning systems used to detect incoming nuclear weapons. NORAD reports indicate that a capability called 'the Pathfinder' fuses data from the military, commercial, and government sensors to create a common operating picture. Previously, data stayed in separate systems, preventing NORAD from seeing the whole picture. Pathfinder, however, can fuse data from more than 300 sensors to build its common operating picture. This has enabled operators to assess threats faster.²⁰

Expert Systems: In the AI context, expert systems are computer programs designed to mimic the decision-making ability of human experts in specific domains, using knowledge representation and reasoning mechanisms to solve complex problems or provide recommendations. Expert knowledge systems use machine learning and artificial intelligence to simulate the behavior or judgment of domain experts and these systems are able to improve their performance over time as they gain more experience, just as humans do.²¹ Expert systems accumulate experience and facts in a knowledge base, and are able to integrate them with an inference or rules engine, defined as a set of rules for applying the knowledge base to situations provided to the program.^{22,23} In NC3 environments, expert systems enhance situational awareness, optimize

19 U.S. Army CCDC Army Research Laboratory Public Affairs, "Machine learning algorithms promise better situational awareness," June 22, 2020, https://www.army.mil/article/236647/machine_learning_algorithms_promise_better_situational_awareness

20 Nathan Strout, "NORAD is Using Artificial Intelligence to See the Threats that it Used to Miss," *C4ISRNET*, March 1, 2021, <https://www.c4isrnet.com/artificial-intelligence/2021/03/01/norad-is-using-artificial-intelligence-to-see-the-threats-it-used-to-miss/>.

21 Nihad Hassan, "AI vs. machine learning vs. deep learning: Key differences," *TechTarget*, August 19, 2025, <https://www.techtarget.com/searchenterpriseai/tip/AI-vs-machine-learning-vs-deep-learning-Key-differences>.

22 Alexander S. Gillis and John Moore, "What are knowledge-based systems (KBSeS)?" *TechTarget*, January 27, 2025, <https://www.techtarget.com/searchcio/definition/knowledge-based-systems-KBS>.

23 Ben Lutkevitch, "Definition: Expert System" *TechTarget*, August 26, 2024, <https://www.techtarget.com/searchenterpriseai/definition/expert-system>.

information flow, and strengthen cybersecurity, potentially shortening the decision-making timeline during crisis scenarios. Their deterministic nature makes them particularly suitable for nuclear command contexts where predictability and reliability are paramount concerns. Expert systems deployed within nuclear command infrastructures showcase performance improvements and ability to reduce decision timelines.

Computer Vision Systems: Computer vision is the field of computer science that enables computers to identify, track, and understand objects and people in images and videos. Computer vision seeks to perform and automate tasks that replicate human capabilities, such as ways that humans see, process, analyze, and interpret. Computer vision systems utilize machine learning to analyze images and videos and are able to execute tasks like object recognition and image classification.²⁴ Within the NC3 context, computer vision systems provide real-time monitoring of mobile launch platforms or nuclear submarines. These systems employ satellite imagery and other advanced imaging technologies to identify weapon types and deployment patterns. Computer vision systems utilize machine learning to analyze images and videos, enabling tasks like object recognition and image classification. Analysts leverage these systems to capture images in order to inform on the nuclear activities of nuclear-armed states.²⁵

Natural Language Processing: Natural language processing (NLP) is a subfield of computer science and AI that uses machine learning to enable computers to understand and communicate with human beings. NLP enables computers and digital devices to recognize, understand, and generate text and speech by combining computational linguistics. It is a rule-based modeling of human language, together with statistical modeling, machine learning, and deep learning models.²⁶ Within the NC3 context, NLP systems could be and often are applied to analyze communications for relevant information and assist in crisis communications. These systems can monitor and analyze diplomatic messages, intelligence reports, and other textual data for indicators of escalating tensions or strategic shifts. Modern NLP systems are able to identify subtle changes in communication patterns that might indicate shifts in nuclear posture or alert status. NLP can be leveraged for rapid threat assessment, communications analysis, and early warning systems.²⁷

24 “What is Computer Vision?” Azure, Microsoft, accessed April 2025, <https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-computer-vision#object-classification>.

25 “Submarine Detection and Monitoring: Open Source Tools and Technologies,” Nuclear Threat Initiative, October 17, 2024, <https://www.nti.org/analysis/articles/submarine-detection-and-monitoring-open-source-tools-and-technologies/>.

26 Cole Stryker and Jim Holdsworth, “What is NLP (Natural Language Processing)?” IBM, August 11, 2024, <https://www.ibm.com/think/topics/natural-language-processing>.

27 Haleema Saadia, Tynchtykbek Israilov, Ekaterina Mikhalevich, and Jonas Sandbrink, “AI and Nuclear Decisions: Toward an Arms Control Framework,” *Contemporary Security Policy*, March 8, 2025, <https://www.tandfonline.com/doi/full/10.1080/13523260.2025.2474869>.

Predictive Analytics: Predictive analytics is an advanced form of data analytics that attempts to predict what happens next. Based on historical data and current inputs, statistical models assess probabilities of various scenarios. Advancements in big data machine learning have also helped expand predictive analytics capabilities.²⁸ Within the NC3 context, predictive analytics can analyze patterns of nuclear weapons deployments, force direction and dispositions, and other indicators to provide early warning of potential nuclear threats or escalation risks. Predictive analytics platforms integrate multiple data streams to generate comprehensive threat assessments with confidence intervals and uncertainty quantification that help commanders make more informed decisions. Scholars indicate that with predictive analytics, AI would be able to pre-emptively analyze troop movements, supply lines, and other intelligence to predict threats before they emerge.²⁹

Potential Limitations and Challenges

Several factors could delay or limit the trajectory of AI systems:

- » **Technical Challenges:** Reliable adversarial robustness, safe autonomous operation in degraded communications environments, and truly trustworthy reasoning remain difficult technical problems that may require more time to solve adequately for military contexts.
- » **Institutional Barriers:** Military organizations may struggle with integration, requiring significant changes to doctrine, training, and organizational structures that typically evolve slowly even with strong leadership support.
- » **Ethical and Legal Frameworks:** The development of comprehensive governance frameworks addressing AI use in conflict—particularly for systems with increasing autonomy—could impose necessary but time-consuming constraints on deployment.
- » **Data Limitations:** Military-specific training data remains limited compared to commercial datasets, potentially slowing the development of specialized military applications.
- » **Cybersecurity Vulnerabilities:** As AI systems become more central to military operations, they present attractive targets for adversaries, potentially limiting deployment until robust security measures are developed.

28 Catherine Cote, “What is Predictive Analysis? 5 examples,” Harvard Business School Online, October 26, 2001, <https://online.hbs.edu/blog/post/predictive-analytics>.

29 Alice Saltini, “AI and Nuclear Command, Control, and Communications: P5 Perspectives,” European Leadership Network, November 2023, https://europeanleadershipnetwork.org/wp-content/uploads/2023/11/AVC-Final-Report_online-version.pdf.

Real World Multi-Agent System Use Cases

Enterprise

- » **Goldman Sachs' AI Research Assistant:** Goldman Sachs developed a multi-agent system where specialized agents handle different aspects of financial research—one extracting data from earnings reports, another analyzing market trends, and a third generating investment recommendations—collaborating to produce comprehensive analyses for investment bankers.³⁰
- » **Walmart's Supply Chain Optimization:** Walmart implemented a multi-agent AI system where different agents monitor inventory levels, predict demand fluctuations, optimize shipping routes, and coordinate with suppliers, resulting in a 17% reduction in stockouts and improved logistics efficiency.³¹
- » **Salesforce's Einstein GPT for Customer Service:** Salesforce deployed a multi-agent system where one agent handles initial customer inquiries, another researches solutions in knowledge bases, and a third drafts personalized responses, with a supervisor agent ensuring quality and consistency across the entire customer interaction.³²

Military

- » **Israel's Lavender System:** The IDF uses this multi-agent intelligence system where different AI components analyze satellite imagery, process signals intelligence, track movement patterns, and identify potential targets, with coordination agents prioritizing information for human decision-makers.³³
- » **Ukraine's Delta Situational Awareness Platform:** Ukraine's military employs a multi-agent system combining drone footage analysis, electronic warfare signal processing, and battlefield mapping agents that collaborate to create real-time operational pictures shared across command units.³⁴
- » **DARPA's OFFSET Program:** The Offensive Swarm-Enabled Tactics program utilizes multiple specialized AI agents controlling drone swarms, where planning agents develop mission strategies, navigation agents optimize paths, and sensory agents gather and interpret battlefield intelligence, all coordinating to accomplish complex missions while maintaining human oversight.³⁵

30 "Goldman Sachs launches AI Agents Worldwide," *Reuters*, June 23, 2025, <https://www.reuters.com/business/goldman-sachs-launches-ai-assistant-firmwide-memo-shows-2025-06-23/>.

31 Hari Vasudev, "Inside Walmart's Strategy for Building an Agentic Future," *Walmart News*, May 29, 2025, <https://corporate.walmart.com/news/2025/05/29/inside-walmarts-strategy-for-building-an-agentic-future>.

32 "Meet Einstein Service Agent: Salesforce's Autonomous AI Agent to Revolutionize Chatbox Experiences," *Salesforce*, July 17, 2024, <https://www.salesforce.com/news/stories/einstein-service-agent-announcement/>.

33 Bethan McKernan and Harry Davies, "The Machine Did it Coldly: Israel Used AI to Identify 37,000 Hamas Targets," *The Guardian*, April 3, 2024, <https://www.theguardian.com/world/2024/apr/03/israel-gaza-ai-database-hamas-airstrikes>.

34 Kateryna Bondar, "Does Ukraine Already Have Functional CJADC2 Technology?," *CSIS*, December 11, 2024, <https://www.csis.org/analysis/does-ukraine-already-have-functional-cjad2-technology>.

35 "DARPA Offset Second Swarm Sprint Pursuing State-of-the-Art Solutions," *Defense Systems Information Analysis Center*, May 30, 2018, <https://dsiac.dtic.mil/articles/darpa-offset-second-swarm-sprint-pursuing-state-of-the-art-solutions/>.

Glossary of Terms

AI agent

An autonomous or semi-autonomous AI system designed to perceive its environment, make decisions, and take actions to achieve specific goals. Agents can operate independently or semi-independently within defined parameters, often utilizing planning and problem-solving capabilities. These agents can operate independently or as part of a larger system, making decisions based on input data and predefined objectives, and are often used to automate routine tasks or provide interactive services.

Multi-agent system

A network of multiple AI agents that collaborate, communicate, and coordinate to solve complex problems. These systems distribute tasks across specialized agents, creating emergent capabilities beyond what any single agent could accomplish.

Scaling laws

Mathematical relationships showing how AI performance improves predictably with increased resources (model size, dataset size, compute).

Parameter counts

The number of adjustable values (weights and biases) in an AI model. Modern large language models contain billions, or even trillions of these values, with higher counts generally enabling more sophisticated capabilities, all else equal.

Multimodal understanding

The ability of AI systems to process and interpret multiple forms of information (text, images, audio, video) cohesively. This enables more comprehensive analysis of complex situations by integrating diverse data types.

Chain of Thought Reasoning

Step-by-step logical processes that AI systems use to arrive at conclusions or decisions. These chains connect premises to conclusions through explicit reasoning steps, making the AI's thinking process more transparent and verifiable.

Test-time scaling

Improving model performance by using more computation when running inference, like implementing step-by-step reasoning, without increasing model size.

Sparse activation patterns

Neural network designs where only a small fraction of neurons activate for any given input, creating efficiency by using specialized pathways rather than the entire network.

Maintaining context	The ability of AI systems to retain and utilize relevant information over extended interactions or complex workflows. This enables consistent understanding of ongoing situations and prevents the loss of critical details during extended operations.
Retrieval-augmented generation	Techniques that enhance AI outputs by combining language model generation with information retrieved from external sources like databases or documents, reducing hallucinations and providing access to more current information.
Context window constraints	Fixed limits on how much context an AI model can process at once, typically measured in tokens. These constraints affect a model's ability to maintain coherence across long inputs, with recent innovations extending these windows significantly.
Longer context windows	The expanded capacity of AI models to retain and process larger amounts of information in a single interaction. Modern systems can now consider thousands to millions of tokens of context, allowing them to analyze entire documents, conversations, or data streams without losing earlier information.
Better tool integration	The improved capability of AI systems to seamlessly utilize and connect with external software, APIs, and specialized tools. This allows AI systems to extend their capabilities beyond their core functions by leveraging other purpose-built systems.

INSTITUTE FOR SECURITY AND TECHNOLOGY

www.securityandtechnology.org

info@securityandtechnology.org

Copyright 2025, The Institute for Security and Technology