AUGUST 2026

SPECIAL COMPETITIVE
STUDIES PROJECT

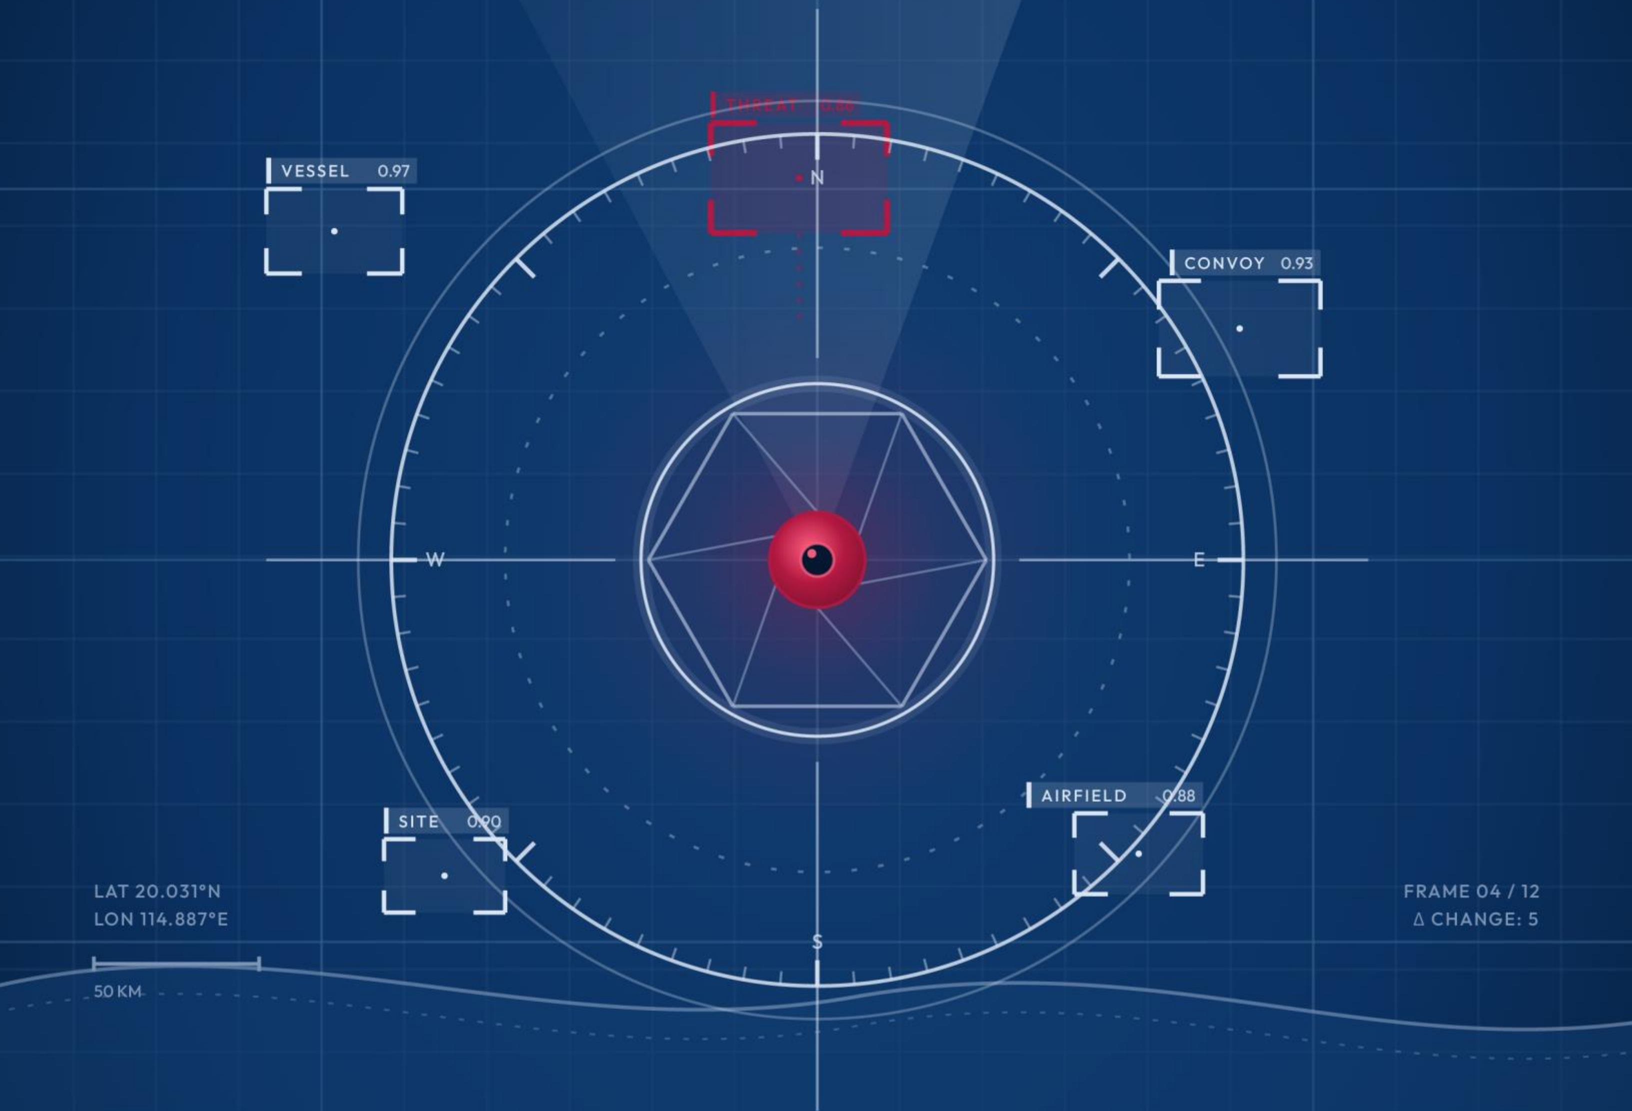

Sensing the Threat:
Operationalizing Computer Vision for Intelligence and Warfighting

Sensing the Threat:

Operationalizing Computer Vision for Intelligence and Warfighting

By Shelby Pierson

Former Director for Analysis, National Geospatial-Intelligence Agency

Disclaimer: This paper is the author's personal work and does not represent the views of the Special Competitive Studies Project, the U.S. Government, the Intelligence Community, the Central Intelligence Agency, the National Geospatial-Intelligence Agency, or any other agency or department. AI-enabled tools were used for limited editorial assistance. All substantive analysis and conclusions are the author's. Due to public prepublication review timelines, the information in this paper is current as of 7 January 2026. Certain figures included in this paper are derived from publicly available materials provided by commercial entities and are used for illustrative purposes only. Their inclusion does not constitute endorsement of any company, product, or service.

Graphics in this report were provided by Vantor and are reproduced with permission.

Executive Summary

In October 1962, photo interpreters at the National Photographic Interpretation Center found Soviet missiles in Cuba and gave an American president thirteen days to act. The lesson of that moment has never changed: the side that sees first, decides first. What has changed is the scale and speed of seeing. Sensors now generate more imagery in a day than the entire Cold War enterprise processed in a year, and no workforce—however skilled—can keep pace with it unaided.

The 2025 National Security Strategy calls on the United States to maintain technological superiority, strengthen AI-enabled military capabilities, and sustain deterrence in an intensifying competition.[1] For more than six decades, the United States has pursued the automation of imagery analysis to increase speed, reduce analytic burden, and preserve decision advantage.[2] The technology has finally caught up to the ambition. The institutions have not. Despite major advances in computer vision, U.S. intelligence and defense organizations continue to rely predominantly on human-centric workflows for imagery exploitation. That gap—between what is technically mature and what is operationally fielded—is no longer an inefficiency. It is a strategic vulnerability.

Computer vision delivers operational value through two proven functions: **change detection,** which identifies deviations from a known baseline across wide areas, and **persistent object monitoring,** which tracks vehicles, vessels, launchers, and other objects across space and time. The technology is mature. The problem is not innovation. It is adoption.

The objectives of the 2025 NSS cannot be met if computer vision remains confined to pilots, demonstrations, and narrow mission threads. Speed and scale now define modern intelligence and military competition.[3] In a future conflict, advantage will accrue to the side that detects change, recognizes patterns, and cues decisions at machine speed. This is not a hypothetical. Peer adversaries are integrating AI-enabled sensing directly into command and control, compressing

[1] 2025 National Security Strategy of the United States of America, The White House (2025).
[2] Central Intelligence Agency Historical Staff, National Photographic Interpretation Center, Volume One: Antecedents and Early Years, 1952–56, NPIC-2 at ii–iii, 1–4 (1972).
[3] 2025 National Security Strategy of the United States of America, The White House (2025).

decision timelines and shedding their dependence on manual analysis.[4] Consider the practical consequence: People's Liberation Army forces detecting force movement, infrastructure change, and logistics activity in near-real time across a theater, while American decision-makers wait on sequential human review and delayed fusion—making choices with less time, less context, and fewer options. History is unsparing toward nations that mistake technological leadership for operational advantage. They are not the same thing, and the difference is decided by institutions, not laboratories.

The 2026–2030 period offers a narrowing window to close this gap. Advances in foundation models, multimodal fusion, edge computing, and on-orbit processing now align with policy momentum and operational necessity.[5] The Department of War's January 2026 AI Acceleration Strategy frames speed, operational integration, and accountable execution as decisive factors.[6] This is the same clock SCSP has been pointing to since its founding: a decisive decade in which the choices made before 2030 determine whether America is positioned to win the techno-economic competition.[7] The United States retains the technological lead. Whether it converts that lead into durable operational advantage is a choice—and the window in which to make it is closing.

---

[4]Gregory C. Allen, Understanding China's AI Strategy: Clues to Chinese Strategic Thinking on Artificial Intelligence and National Security, Center for a New American Security (2019); U.S.–China Economic and Security Review Commission, Annual Report to Congress (2025).
[5]Fei-Fei Li, From Words to Worlds: Spatial Intelligence Is the Next Frontier of AI, Substack (2025).
[6]Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026).
[7]Mid-Decade Challenges to National Competitiveness, Special Competitive Studies Project (2022).

Introduction

The ambition to automate imagery analysis is as old as overhead reconnaissance itself. During the Cold War, the National Photographic Interpretation Center explored early computational methods to help analysts drowning in satellite film.[8] Through the 1980s and 1990s, automatic target recognition and feature-based algorithms delivered incremental gains but remained brittle—sensitive to weather, lighting, and terrain—reinforcing a cultural conviction that only human analysts could reliably interpret complex scenes.

The 2000s brought better sensors, more compute, and cheaper storage, but no fundamental shift in tradecraft. Organizations optimized the "single pane of glass" paradigm, consolidating feeds into two-dimensional dashboards.[9] Meanwhile, the human craft of imagery analysis continued to deliver at the highest levels of consequence—from the identification of the Bin Laden compound to arms transfer monitoring, ballistic missile warning, and battlefield assessment. What never occurred was the parallel effort to mature that craft into model development and enterprise data ontology. The tradecraft advanced. The architecture stood still.

Project Maven was the most significant American effort to break that pattern. Launched in 2017 by then-Deputy Secretary of Defense Bob Work as the Algorithmic Warfare Cross-Functional Team, Maven attacked a concrete operational problem: unmanned aerial systems were generating full-motion video faster than human analysts could possibly process it. Maven's core function was object detection—the automated identification and labeling of objects of interest within imagery and video. By applying deep learning to full-motion video, Maven automated the

---

[8]Central Intelligence Agency Historical Staff, National Photographic Interpretation Center, Volume One: Antecedents and Early Years, 1952–56, NPIC-2 at ii–iii, 1–4 (1972). This paper focuses on geospatial intelligence (GEOINT) analysis and production rather than all-source analysis writ large, examining how computer vision can most immediately transform the geospatial foundation upon which many all-source assessments depend. With respect to sensing modalities, the focus is on computer vision applied to airborne and space-based imagery, including full-motion video from air-breathing platforms and electro-optical data from on-orbit sensors. It does not address facial recognition, ubiquitous terrestrial surveillance, or population-scale monitoring.

[9]Mark Munsell, It's Time to Break Through the Single Pane of Glass, LinkedIn (2024); Brandi Vincent, DOD Wants More than $2B in Fiscal 2027 to Move Beyond 'Fragmented' CJADC2 Deployments, DefenseScoop (2026).

detection of vehicles, equipment, and other tactical objects, sharply reducing manual scanning and accelerating analytic workflows.[10]

Maven succeeded where four decades of predecessors had struggled because it was operationally grounded, tightly scoped, and embedded in real mission workflows. It did not attempt to automate imagery analysis wholesale; it took one high-value function and drove it into existing processes. Analysts moved from exhaustive manual scanning to supervisory, decision-focused roles. Engineers, operators, and analytic leadership iterated models against live mission data. And as contemporary reporting documented, Maven's trajectory was shaped as much by institutional resistance and cultural debate as by technical performance. It succeeded not by bypassing institutional constraints but by navigating them—reframing AI as decision support aligned with existing accountability structures.[11]

The Army's Maven Smart System extended the model, demonstrating how object detection could be sustained through continuous integration, DevSecOps practices, and close collaboration among operators, engineers, and commanders. By the early 2020s, multimodal foundation models and spatial intelligence architectures emerged, enabling systems to reason about three-dimensional environments. Exercises such as Scarlet Dragon showed that automated object detection could dramatically compress targeting timelines and reduce personnel requirements while maintaining compliance with legal and policy constraints. The war in Ukraine demonstrated beyond question that computer vision and commercial imagery could operate at meaningful scale, publicly tracking Russian force movements and infrastructure damage. Yet demonstration has not always translated into sustained institutional adoption. Today, computer vision runs global logistics, infrastructure monitoring, agriculture analytics, urban mapping, and autonomous navigation at planetary scale. The question is no longer whether the technology works. It is why the national security enterprise still treats it as optional.

[10] Katrina Manson, AI Warfare Is Already Here, Bloomberg (2024).
[11] Katrina Manson, AI Warfare Is Already Here, Bloomberg (2024).

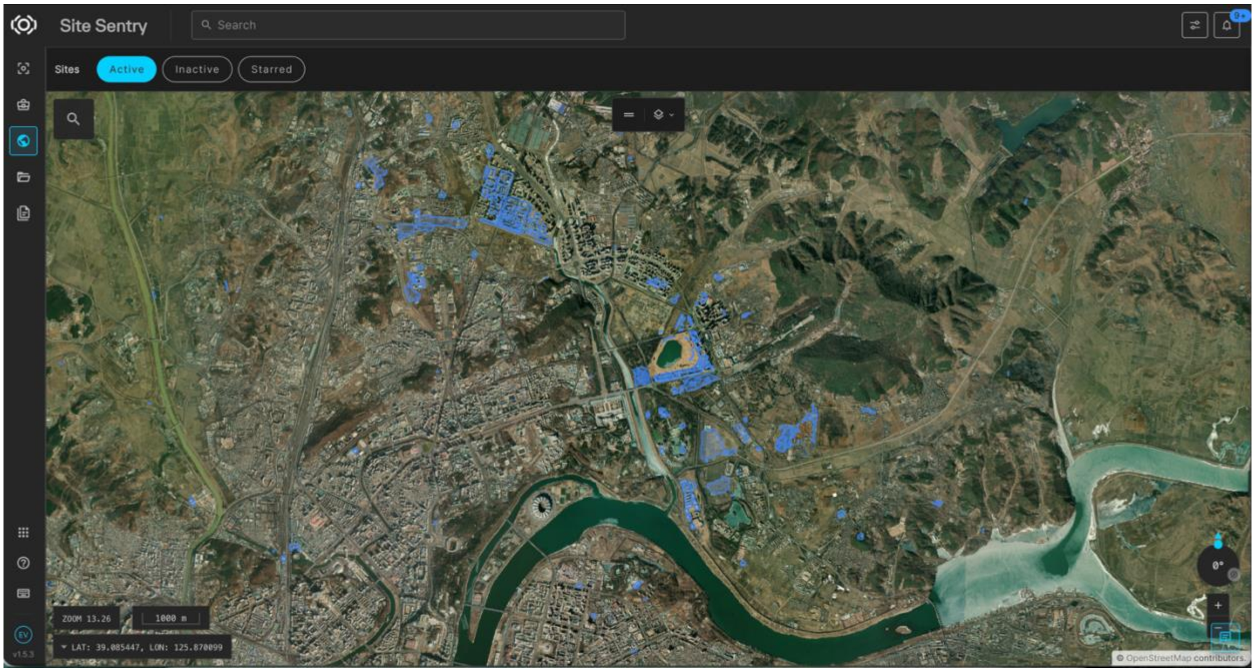

[Figure 1. Automated Change Detection — Computer vision automatically identifies deviations from established baselines across wide geographic areas, enabling analysts to rapidly detect emerging activity without manually scanning entire scenes. By directing attention to meaningful change, automated change detection accelerates warning, supports persistent monitoring, and improves decision speed.]

Barriers to Wider Adoption

Despite Maven's proven capability and the rapid advance of generative, multimodal AI, U.S. intelligence and defense organizations continue to rely heavily on manual imagery exploitation. Analysts still scan imagery, identify objects of interest, and assess change by hand—even as data volume, variety, and velocity grow exponentially. The primary barriers are not technical. They are institutional, cultural, and procedural, and together they inhibit scale.[12]

Data constraints and classification remain the most persistent challenge. Computer vision systems require large volumes of high-quality labeled data, including negative examples and edge cases. In national security contexts, mission-relevant imagery is fragmented across classification levels, compartments, and organizations. Cross-domain transfer processes slow iteration, and classification restrictions routinely prevent developers from accessing the data needed to retrain models as adversaries adapt. These constraints bite hardest in targeting-grade identification, where fine-grained labels must be continuously updated to reflect new variants, camouflage techniques, and operational environments.

Within the National Geospatial-Intelligence Agency and other Intelligence Community elements, data suitability is not simply a matter of volume. Imagery used for analytic production must meet standards for provenance, consistency, and reliability. Labels are not created casually, and negative examples often matter as much as positive detections. For targeting-grade use cases, model outputs must be validated across sensors, conditions, and time. These requirements dramatically raise the bar for training data and slow adoption relative to commercial environments where iterative error is acceptable.[13]

Workflow and cultural resistance further limit adoption. Imagery analysis remains rooted in artisan tradecraft. Analysts are trained to interpret images manually and validate conclusions personally. Automation is treated as advisory rather than authoritative—even when models outperform human-only workflows. Two-dimensional dashboards reinforce mental models that

---

[12] Jennifer Ewbank, The Role of Artificial Intelligence in the U.S. Intelligence Community: Current Uses and Future Developments, Aspen Institute (2024).
[13] Office of the Director of National Intelligence, Intelligence Community Directive 203: Analytic Standards (2015).

privilege static views over temporal and spatial reasoning. The result: computer vision gets bolted onto existing processes rather than used to redesign workflows around machine-speed analysis. Bolting the future onto the past is not modernization. It is a way of postponing it.

Human capital dynamics sit at the center of this resistance. Analysts are trained, evaluated, and celebrated as accountable producers of exceptional judgment, not as supervisors of automated systems. When models outperform humans on detection or classification, the obstacle is not technical disbelief but professional displacement anxiety: where does analytic value reside if machines see first? Absent a clear, practical distinction between machine-enabled detection and human judgment, incentives naturally favor manual validation and conservative use. This is not a failure of workforce adaptation. It is the predictable outcome of accountability structures that have not been updated for human–machine teaming.

Medical radiology offers the instructive parallel. AI has demonstrated strong performance detecting tumors, fractures, and other anomalies in medical imagery. The most effective applications have not replaced clinicians; they have reshaped workflows so machines perform first-pass detection while humans retain responsibility for diagnosis, judgment, and patient outcomes. Adoption has been uneven not because the technology failed, but because institutions struggled to realign incentives, liability frameworks, and professional norms around human–machine collaboration.[14] Imagery analysis faces the same structural challenge. Computer vision excels at detection and monitoring at scale; its value depends on how institutions redefine the human role in interpretation, accountability, and decision-making. This conclusion is reinforced by recent work from both SCSP and the Australian Strategic Policy Institute (ASPI), which similarly argue that the future of intelligence lies in high-performing human-machine teams. Machines should perform the labor of detection, monitoring, and synthesis at scale, while human analysts remain responsible for judgment, context, and accountability.[15]

Cultural resistance extends to commercial and open-source information. Despite the growing maturity of commercial imagery, publicly available information, and commercially derived

[14] Te-Ping Chen & Chao Deng, [Hospitals Are a Proving Ground for What AI Can Do, and What It Can't](), The Wall Street Journal (2026).
[15] [The Future of Intelligence Analysis: AI and Human-Machine Teaming](), Special Competitive Studies Project & Australian Strategic Policy Institute (2024).

analytics, an institutional bias persists toward treating these sources as supplementary rather than foundational. This forfeits the ability to use computer vision-enabled commercial and open-source data as a first analytic layer for trend detection, warning, and cueing. Reframing commercial and open sources as the first tranche of analytic input—fused subsequently with classified intelligence for verification and context—would both normalize computer vision adoption and reduce analytic latency across the community. SCSP's Intelligence Panel has made the underlying case at the enterprise level: the Intelligence Community now operates in an era defined by data that is largely open or commercially available, and open-source assessments increasingly rival what classified production delivers. Treating these sources as foundational is not a concession. It is the logical response to where the data actually lives.[16]

Trust, accountability, and governance gaps compound the problem. Accountability structures were designed for human reasoning, not algorithmic outputs. There is limited doctrine for handling confidence scores, uncertainty, auditability, or human–machine collaboration in operational contexts.[17]

[16] Intelligence Innovation: Repositioning for Future Technology Competition, Special Competitive Studies Project (2024); Intelligence in an Age of Data-Driven Competition, Special Competitive Studies Project (2022).
[17] Office of the Director of National Intelligence, Intelligence Community Directive 203: Analytic Standards (2015).

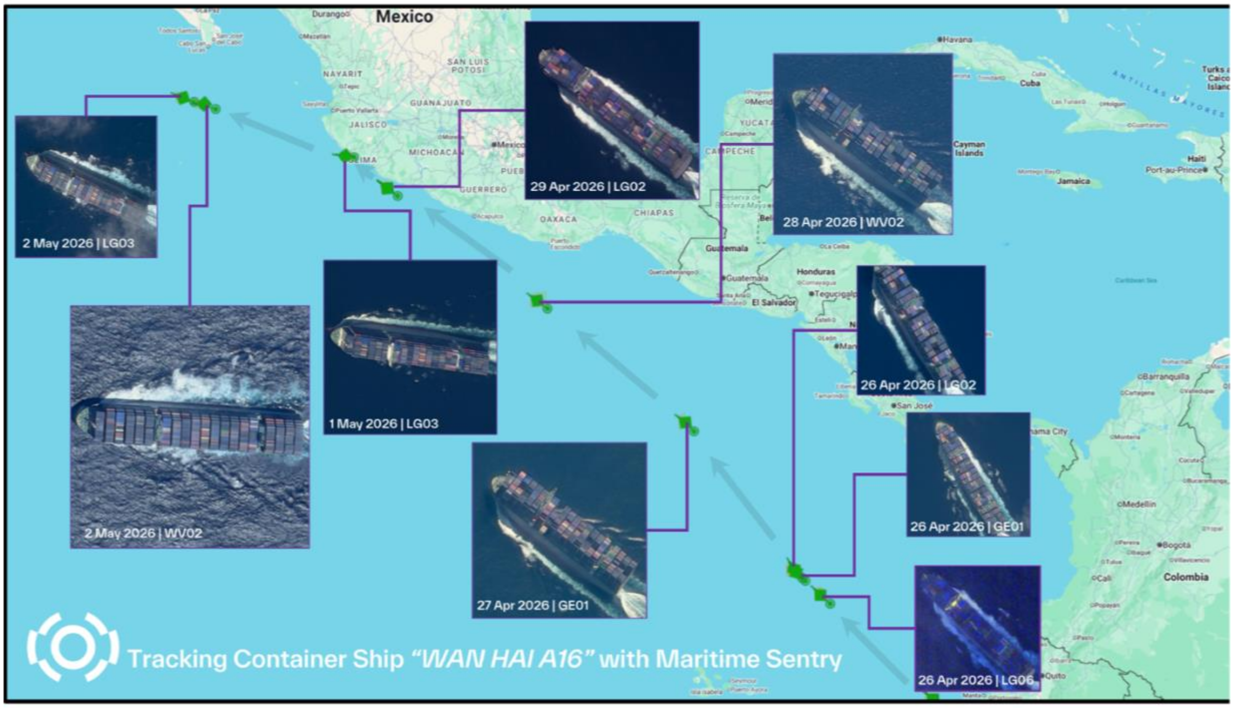

[Figure 2. Persistent Vessel Detection Supporting Maritime Domain Awareness — An automated tipping and cueing engine enabled continuous monitoring of a container ship after it was identified as a vessel of interest based on indicators of potentially suspicious activity. Persistent object monitoring allows analysts to detect anomalous behavior, establish patterns of movement, and maintain continuous maritime situational awareness while reducing reliance on manual observation.]

Without governance frameworks that establish justified confidence in system performance, organizations default to manual review and restrict AI's operational role. These challenges are magnified for targeting applications, where legal, ethical, and escalation considerations are paramount.[18] Analytic accountability within the IC is both a strength and a constraint. Analysts are personally responsible for judgments that may influence military action or policy decisions. As a result, AI systems are often held to a higher standard than human analysts—expected to be simultaneously accurate, explainable, and robust to adversarial manipulation. Without clear doctrine for how machine confidence, uncertainty, and error propagate through analytic

[18]For analytic purposes, knowing that an object is a tank or truck is often sufficient. Targeting demands greater specificity—distinguishing a T-72 from a T-90 carries operational, legal, and escalation implications. Fine-grained targeting requires high-quality labeled data that is difficult to obtain and maintain as adversaries modify equipment and employ deception. Models trained on static datasets risk rapid obsolescence, requiring continuous data refresh and retraining pipelines.

products, organizations default to conservative use. This is not resistance to technology. It is risk management embedded in tradecraft—and it will not change until the doctrine does.

Acquisition misalignment locks the pattern in place. Traditional acquisition pathways treat AI as static software rather than a continuously learning capability. Approval processes freeze models in time, preventing rapid updates as conditions change. Pilot programs proliferate because there are few clear transition paths to sustained deployment. The result is a familiar cycle: repeated demonstrations of capability without enterprise-scale adoption.

Institutional fragmentation completes the picture. No single authority is responsible for integrating computer vision across intelligence, surveillance, and reconnaissance missions. Services, agencies, and combatant commands pursue parallel efforts, producing duplication and incompatible architectures. Even successful programs struggle to scale beyond their initial context.[19] The Department's recent move to establish an Enterprise Command and Control Program Office reflects growing recognition that fragmented analytics and stove-piped decision support are no longer tenable. Senior leadership has been tasked to consolidate authorities, infrastructure, and support activities across intelligence and operational elements to deliver integrated, AI-enabled decision support at scale.[20] The implication is clear: capabilities such as computer vision cannot be treated as standalone tools or single-program successes. They must be embedded within enterprise C2 architectures to meaningfully affect operational tempo and decision advantage.

Finally, conceptual blind spots persist. Many systems still treat imagery as flat pictures rather than representations of a three-dimensional physical environment—limiting spatial reasoning, persistence, and prediction. Until workflows and tools are designed around the physical reality of the battlespace, adoption will remain partial. In complex terrain and border environments, episodic two-dimensional analysis is particularly ill-suited to detecting persistent activity, reinforcing the need for spatially aware, machine-enabled monitoring.

---

[19] Brandi Vincent, DOD Maps Out Plan for New Enterprise Command and Control Program Office, C2 Suite, DefenseScoop (2026).
[20] Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026); Brandi Vincent, DOD Maps Out Plan for New Enterprise Command and Control Program Office, C2 Suite, DefenseScoop (2026).

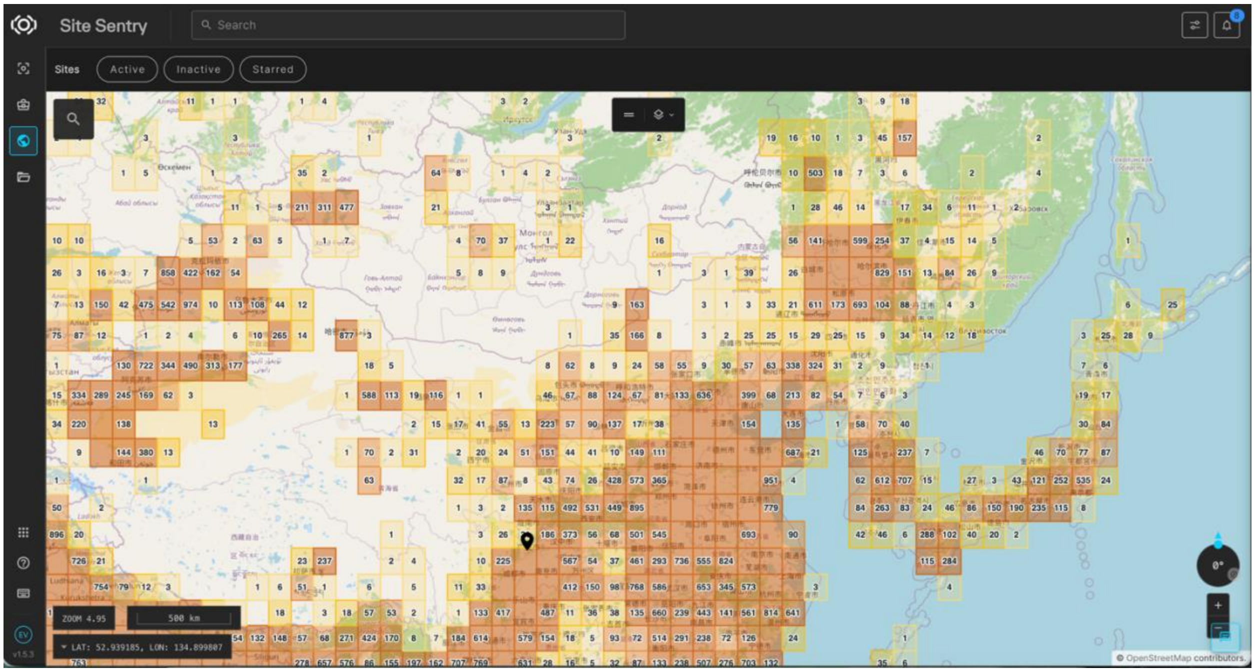

[Figure 3. Computer Vision–Enabled Detection in Complex Terrain — Computer vision automatically identifies and prioritizes sites of potential interest across dense urban terrain, allowing analysts to focus on locations requiring further assessment rather than manually reviewing every image. Automated detection improves the speed, consistency, and scale of intelligence monitoring while enabling human analysts to concentrate on interpretation and operational decision-making.]

How AI-for-GEOINT Is Used in China and Israel

The cost of delayed adoption is best understood by comparison. China and Israel offer two distinct but instructive models of operationalized computer vision—each illuminating a different dimension of the challenge facing the United States.

China: Scale, Diffusion, and Automation First

China's approach emphasizes scale and diffusion over isolated technical superiority. Chinese military and security organizations have integrated computer vision across a wide range of missions—maritime domain awareness, border surveillance, infrastructure monitoring, and internal security.[21] These systems automate change detection and persistent object monitoring across large geographic areas, reducing reliance on human analysts and compressing decision timelines. In the maritime domain, computer vision supports persistent tracking of commercial and military vessels, flagging anomalous behavior and deviations from baseline patterns. Automated change detection identifies new construction, equipment emplacement, and logistics activity in contested and sensitive areas. Critically, these capabilities are integrated into command-and-control architectures rather than treated as standalone analytic tools.

The defining feature of China's approach is diffusion. Computer vision is deployed across multiple echelons and mission areas, not confined to pilots or elite units.[22] This reflects a deliberate strategy to institutionalize AI-enabled sensing as a routine component of operations. Individual Chinese systems may not always match their American counterparts in sophistication. Yet their integration into operational workflows creates cumulative advantage. Their integration into operational workflows creates cumulative advantage. Open-source reporting on PLA maritime and space-based ISR shows continuous automated tracking of vessels, construction, and logistics across contested regions—persistence and scale, not perfection, generating decision-speed advantage.

[21] Jacob Stokes, Military Artificial Intelligence, the People's Liberation Army, and U.S.-China Strategic Competition, Center for a New American Security (2024).
[22] Artificial Eyes: Generative AI in China's Military Intelligence, Recorded Future / Insikt Group (2025); China's Growing Use of Artificial Intelligence in Military Operations, The China Briefing (2025); Memo to the President and Congress on the Adoption of Generative Artificial Intelligence for Military Purposes, in Generative AI: The Future of Innovation Power, Special Competitive Studies Project (2023).

Over the past several years, China has operationalized computer vision to support persistent maritime domain awareness across the South China Sea. Open-source reporting and commercial imagery analysis show automated vessel tracking integrated with satellite imagery, AIS data, and coastal sensors to identify anomalous behavior, shifts in vessel patterns, and new infrastructure activity around disputed features. Rather than relying on episodic human review, these systems continuously establish baselines and flag deviations—unusual loitering, formation changes, construction activity—at scale. The operational advantage does not stem from perfectly classifying every vessel. It stems from persistent, automated change detection that compresses warning timelines and enables earlier cueing of command-and-control decisions.

China's governance model emphasizes rapid operational fielding and iterative deployment, reflecting centralized civil-military integration and state-directed technology development. U.S. adoption, by comparison, proceeds through more extensive acquisition, testing, legal review, and oversight.[23] China's model fields automated monitoring and targeting-relevant analytics faster, even at the cost of transparency and interpretability. The strategic logic is unambiguous: speed, scale, and persistence.

Israel: Deep Operational Integration Under Pressure

Israel offers a different model, centered on depth of integration rather than geographic scale. Confronting dense sensor environments, continuous threats, and limited analytic manpower, Israel has treated automation as a necessity rather than an enhancement. Recent Israeli operations—including highly asymmetric actions against Hezbollah infrastructure and networks—illustrate an operational culture that fuses rapid detection, automated cueing, and tightly coupled targeting decisions: deep institutional reliance on machine-assisted tempo rather than episodic automation.[24]

[23] Final Report, National Security Commission on Artificial Intelligence (2021).
[24] Katrina Manson, AI Warfare Is Already Here, Bloomberg (2024); Kanishka Singh, U.S. Looking at Report That Israel Used AI to Identify Bombing Targets in Gaza, Reuters (2024); Geoff Brumfiel, Israel Is Using an AI System to Find Targets in Gaza. Experts Say It's Just the Start, National Public Radio (2023).

Recent Israeli military operations illustrate how computer vision integrates into monitoring and targeting workflows.[25] Open-source reporting indicates that automated detection and change monitoring cued analysts and operators to emerging activity—movement near known infrastructure, changes at suspected facilities, deviations from established patterns—allowing human teams to focus on validation, context, and decision-making rather than manual scanning. This human–machine teaming model enabled continuous operations under intense tempo, where delays of hours rather than minutes carried operational consequences. Computer vision functioned as a first-pass sensing layer, not a replacement for judgment—demonstrating how automation becomes indispensable once operational demand exceeds human bandwidth.

In Israeli operations, computer vision is integrated directly into monitoring, target development, and operational planning.[26] Change detection identifies new activity in highly dynamic environments. Persistent object monitoring tracks individuals, vehicles, and infrastructure over time. These capabilities are coupled tightly to decision-making rather than isolated in analytic stovepipes.

Israel's experience demonstrates a hard truth: computer vision becomes indispensable when data volumes exceed human capacity. Automation is not deployed to replace analysts but to allow analysts and operators to function at operational tempo. Israel's approach also exposes governance tradeoffs. Rapid deployment and integration have generated controversy and debate around accountability and civilian harm—underscoring the necessity of governance frameworks that enable automation while maintaining legal and ethical controls. But the debates also confirm the underlying dynamic: once operational demand crosses a threshold, automation stops being a choice.

---

[25]Harry Davies, et al., 'The Gospel': How Israel Uses AI to Select Bombing Targets in Gaza, The Guardian (2023).

[26]National Geospatial-Intelligence Agency, National Geospatial-Intelligence Agency Capabilities and Contributions (2025).

Expanding Beyond Maven

The period between 2026 and 2030 represents a narrowing window in which the United States can convert its advances in computer vision into durable operational advantage. The window exists not because the technology is new, but because the next generation of intelligence architectures, satellite constellations, command-and-control systems, and AI-enabled workflows is being designed and fielded now. Decisions made in this period will shape operational capability for a generation. Integration deferred is integration made harder, costlier, and more disruptive. Advances in foundation models have already reduced dependence on large, bespoke labeled datasets, allowing systems to generalize across sensors, environments, and mission contexts with fewer manual annotations.[27]

The United States will capture this innovation only if it adapts its acquisition, integration, and sustainment models accordingly. Absent that adaptation, commercial progress will continue to outpace national security adoption—widening the gap between what is technically possible and what is operationally fielded.[28] Most critically, operationalizing computer vision at scale requires a shift in where processing occurs. Ground-based architectures are no longer sufficient for automated target recognition, persistent monitoring, or time-sensitive decision support. Ground-only models impose latency, bandwidth constraints, and single points of failure incompatible with large sensor constellations and contested environments. Continuing to architect computer vision around ground processing will lock the United States into structural disadvantage as data volumes grow and adversary tempo accelerates.

---

[27] Fei-Fei Li, From Words to Worlds: Spatial Intelligence Is the Next Frontier of AI, Substack (2025); Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026).

[28] Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026); DoD Chief Digital and AI Office Introduces New Advantage at Industry Day, U.S. Department of War (2024).

On-orbit computer vision should therefore be treated as a strategic requirement, not an experimental capability. Processing at the sensor layer enables earlier detection, prioritization, and filtering; reduces unnecessary downlink; and sustains persistence when communications are degraded or denied.[29] For targeting-related missions, on-orbit processing generates time-sensitive insight at machine speed rather than at the pace of human-centered, ground-based pipelines. This shift is foundational to competing against adversaries explicitly designing to compress decision timelines.

Nor does this argument stand alone. SCSP's Offset-X strategy places sensing—alongside AI and autonomy—at the foundation of a new Joint Warfighting Concept: a triad that, adopted as an integrated system, delivers informational, decisional, and lethality advantage.[30] Computer vision is the connective tissue between the sensing layer and the decision layer of that triad. Without operationalizing computer vision, investments in sensing, AI, and autonomy cannot fully translate into decision advantage. The value of increasingly capable sensors depends on the ability to rapidly detect, prioritize, and deliver meaningful information to decision-makers at operational tempo. Advances in edge compute and space-qualified AI chips make the transition technically feasible within the current planning horizon. What remains is not a technology gap. It is a policy and governance decision. To succeed, the United States must pair on-orbit processing with governance frameworks that support continuous model updates, validation, and retraining while preserving analytic accountability and legal oversight.[31] Computer vision must be treated as a living capability, not a static system. The choice before U.S. leaders is not whether computer vision will mature—it will. The choice is whether the national security enterprise reorganizes itself to exploit it. The 2026–2030 period offers a rare alignment of technology, acquisition, and governance around operational use. This period also coincides with a generational transition in

---

[29] Theresa Hitchens, NRO Awards Three New Contracts for Commercial Satellite Data, Breaking Defense (2026); Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026).
[30] Offset-X Evolved: A Technology-Powered Strategy for the U.S. Military, Special Competitive Studies Project (2025); Offset-X: Closing the Deterrence Gap and Building the Future Joint Force, Special Competitive Studies Project (2023).
[31] Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026); Defense Innovation Board, AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense, U.S. Department of Defense (2019).

the analytic workforce, creating a unique opportunity to modernize both technology and tradecraft simultaneously. Failing to act within this window risks hard-coding today's human-bound workflows into tomorrow's architectures—and conceding tempo advantage to competitors already designing for machine-speed operations.

Recommendations: Compete to Win

Recommendations that begin with governance begin in the wrong place. The United States does not suffer from a shortage of oversight. It suffers from a shortage of fielded capability. In 1940, France possessed more tanks than Germany—and in many respects better ones. Germany won the Battle of France in six weeks because it integrated radios, doctrine, and organization around a new way of fighting, while France perfected the defenses of the last war. Innovation did not decide that campaign. Adoption did. The danger for the United States today is the same strategic narcissism that undid France: the assumption that because we invented the technology, we will be the ones to benefit from it. Strategy is implementation. The recommendations below are sequenced accordingly: field first, buy for speed, build the architecture of machine-speed warfare, and hold commanders accountable for outcomes—not process.

Field at the Speed of the Fight

Stop piloting. Start fielding. Maven proved that computer vision succeeds when it is embedded in real operational workflows and iterated against real mission data.[32] That lesson is nearly a decade old, and the enterprise is still relearning it one pilot at a time. Every computer vision effort should now carry three things from inception: an operational sponsor, explicit transition criteria, and a termination date. Capabilities that meet the criteria transition to sustained operations on a mandated timeline. Capabilities that do not are killed, and their resources reallocated. Transition or terminate. Endless experimentation is not caution—it is a decision to lose slowly, made one deferral at a time. Mission-thread-based fielding, executed under operational conditions rather than abstract performance metrics, is also how cultural resistance dies: not by argument, but by demonstrated advantage in the fight.

Make machine-first detection the default—and manual scanning the exception. Computer vision is a decision-acceleration capability. Human effort belongs on judgment, interpretation, and escalation management—not on triage that machines perform faster and more consistently. Workflows should be inverted accordingly: automated first-pass detection and monitoring as the standard operating baseline, with manual scanning reserved for cases that demand it and

[32]Katrina Manson, AI Warfare Is Already Here, Bloomberg (2024).

justified as such. This is not a diminishment of the analyst. It is the restoration of the analyst to the work that actually requires one—meaning, intent, and consequence—consistent with Intelligence Community accountability norms and mirroring the adoption pattern in medical imaging.[33] Machines see first. Humans decide. Any force that arranges it otherwise will find itself reacting to an enemy that did not.

Make Acquisition a Weapon

Buy learning systems, not frozen software. Traditional acquisition freezes AI at the point of approval, guaranteeing obsolescence as adversaries adapt and environments change.[34] For computer vision—where model relevance depends on continuous data refresh—this is not a procedural flaw. It is unilateral disarmament on a schedule. Acquisition authorities should approve computer vision as a living capability, with retraining, updating, and performance monitoring built into the authorization itself, not bolted on through subsequent reviews. And the enterprise should buy commercial first: commercial imagery, commercial analytics, and commercially derived computer vision as the first tranche of analytic input, fused with classified collection for verification and context. The commercial sector will continue advancing regardless of the pace of government adoption.[35]

Build GEOINT-scale training and evaluation infrastructure now. The Intelligence Community and the Department of War should establish continuous labeling, retraining, and validation pipelines, with shared infrastructure for training and testing models across sensors and mission contexts.[36] That infrastructure must emphasize negative examples, edge cases, and cross-sensor validation—the evidentiary standards intelligence production demands. This is how the enterprise scales without surrendering analytic rigor: not by lowering the bar, but by industrializing the

---

[33] The Future of Intelligence Analysis: AI and Human-Machine Teaming, Special Competitive Studies Project & Australian Strategic Policy Institute (2024).
[34] Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026); Defense Innovation Board, AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense, U.S. Department of Defense (2019).
[35] Intelligence Innovation: Repositioning for Future Technology Competition, Special Competitive Studies Project (2024).
[36] Office of the Director of National Intelligence, Intelligence Community Directive 203: Analytic Standards (2015); National Geospatial-Intelligence Agency, National Geospatial-Intelligence Agency Capabilities and Contributions (2025).

means of clearing it. SCSP's Intelligence Panel called for exactly this in its generative AI work: scale cutting-edge AI across the intelligence cycle, and build the data and evaluation backbone that makes scaling responsible.[37]

Build the Architecture of Machine-Speed Warfare

Move processing to the sensor. Ground-centric architectures impose latency, bandwidth constraints, and single points of failure that are incompatible with contested, data-dense environments. On-orbit and edge processing enable detection, prioritization, and filtering at the point of collection—preserving decision advantage even when communications are degraded or denied.[38] The next generation of constellations and command-and-control systems is being designed now; edge processing must be mandated into those designs as a strategic requirement, not appended later as an experimental add-on. Architecting future sensing enterprises around ground-only processing risks embedding latency into systems that must compete at machine speed.

Invest in spatial intelligence and three-dimensional world models. Legacy two-dimensional paradigms limit persistence and spatial reasoning. Adversaries do not operate on flat pictures; they operate in physical terrain, across time. Spatial intelligence architectures allow computer vision to reason across time, space, and modality—closing the conceptual blind spots that keep the technology from delivering its full operational value and aligning analytic tools with the physical reality of the battlespace.

---

[37]Memo to the President and Congress on the Implications of Generative Artificial Intelligence for the U.S. Intelligence Community, in Generative AI: The Future of Innovation Power, Special Competitive Studies Project (2023).
[38]Theresa Hitchens, NRO Awards Three New Contracts for Commercial Satellite Data,, Breaking Defense (2026); Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026).

Command Accountability for Outcomes

Put one leader on the hook—and give that leader the authority to act. Institutional fragmentation persists because no one owns the outcome. Responsibility for AI-enabled ISR is scattered across services, agencies, and program offices, producing duplicated effort, incompatible architectures, and pilots that never transition. The answer is not a new bureaucracy or another coordinating council—it is command accountability. The Secretary of War, in coordination with the Director of National Intelligence, should designate a single accountable official with the authority to set shared performance outcomes, enforce common evaluation benchmarks, force transitions, and kill duplicative programs. This is the execution model the Department's own 2026 AI Acceleration Strategy demands: single accountable leaders, aggressive timelines, and measurable outcomes through its Pace-Setting Projects.[39] Accountability is not oversight layered on top of the mission. It is how militaries win.

Write doctrine that lets the enterprise say yes. Existing accountability structures were built for human judgment, not machine-generated detections, confidence scores, and uncertainty estimates. In the absence of doctrine for how machine confidence and error propagate through analytic and operational products, organizations default to manual review—not from timidity, but because tradecraft demands it. The fix is not to lower standards. It is to write the doctrine: explicit guidance on validation, auditability, and human–machine collaboration that establishes justified confidence and shifts the default from restriction to responsible operational use, while preserving the legal, ethical, and escalation controls that distinguish American power from its competitors. Doctrine that cannot keep pace with the fight does not restrain the enemy. It restrains us.

[39] Artificial Intelligence Strategy for the Department of War: Accelerating America's Military AI Dominance, U.S. Department of War (2026).

Conclusion

Computer vision is no longer experimental or peripheral. It sits at the center of how modern intelligence and warfighting will be conducted as sensor proliferation, data volume, and decision speed outpace human-only workflows. The question is not whether computer vision will be used. It is whether it will be operationalized in a way that preserves analytic rigor, supports responsible targeting, and sustains decision advantage.

The evidence is increasingly clear. Maven demonstrated that computer vision can be operationalized when narrowly scoped and institutionally aligned—but object detection alone cannot carry the full range of intelligence and targeting missions. China is operationalizing at scale to compress decision timelines. Israel has shown that when operational tempo exceeds human capacity, automation becomes unavoidable. The United States retains technical superiority. Superiority unfielded is superiority forfeited. Without institutional alignment and enterprise-scale adoption, the American advantage will erode—not through any single failure, but through the quiet accumulation of pilots that never transitioned and decisions never made. Acting decisively within this window offers the opportunity to translate technological leadership into durable strategic advantage—at the moment when it is both attainable and matters most. The interpreters at NPIC gave the president thirteen days. The next crisis may take thirteen hours. Whether America sees first will be decided now.