Human Factors in Robots, Drones and Unmanned Systems, Vol. 218, 2026, 7–17
https://doi.org/10.54941/ahfe1007673

The Risks, Challenges, and Potential Opportunities With GenAI

Kristin E. Schaefer[1], John F. Tomaselli[2], Larry Parrotte[2], Brandon Taylor[2], Antonio Magana[2], Henry Reimert[1], Selena Hamilton[3], Maggie Wigness[1], and Daniel Cassenti[1]

[1]DEVCOM Army Research Laboratory, Aberdeen Proving Ground, MD 21005, USA
[2]US Army AI Integration Center (AI2C), Pittsburgh, PA 15206, USA
[3]Energetics Technology Center, Indian Head, MD 20640, USA

ABSTRACT

Artificial intelligence (AI) is a field where the masses offer declarations about novel advancements to machine intelligence and the everyday person feels like an AI "expert" in Generative AI (GenAI), such as ChatGPT and DALL-E. While 80+ years of research has led to the potential for GenAI to create new, "original" content, the public ought to understand that GenAI's abilities are predicated on processing massive datasets. These datasets have many potential risks, including overtraining or novel datasets, foundational data science and metadata to AI models to cause incorrect decisions, bypass security, or extract sensitive information. Effective, trusted teaming with AI-agentic teams remains a critical research and development objective. Further, AI effectiveness becomes irrelevant if a human does not understand or trust the AI. This paper provides the foundations of AI and risks of GenAI, followed by a Use Case example of data management from a sensor edge node through actionable intelligence describing AI. This Use Case will walk through a data science strategy underpinning AI for enhancing trusted AI-agentic teaming, outlining the scientific research, challenges, and risks that can occur at each step that can directly impact the trusted relationship. Distribution A: Approved for public release: distribution is unlimited.

Keywords: Artificial intelligence, Human-AI teaming, Live virtual constructive simulation, Test and evaluation

INTRODUCTION

Generative artificial intelligence (GenAI) is described as deep machine learning that is trained to create new or novel outputs (MIT, 2023). Technology is now available to the masses through readily available access to large-language models and chatbots for text-based outputs, image generation and editing, video generation tools, code development, website design, and more. Everywhere you look a new GenAI technology is available, and depending on who is reviewing the area, the list of technologies is ever changing. What does not change is this perception that the everyday person is an "expert" in artificial intelligence (AI). The design of these tools with user friendly front ends and visually appealing outcomes affect a user's trust calibration in the technology, often leading to mistrust, distrust, or over-trust in the outcome. This is in part due to the Dunning-Kruger effect, a cognitive bias where

Received March 21, 2026; Revised April 28, 2026; Accepted May 15, 2026; Available online July 20, 2026

people with limited knowledge or competence in a particular domain greatly overestimate their own knowledge or competence (Kruger & Dunning, 1999). This is common with GenAI technologies where users often confuse using a tool with understanding its underlying technology, also known as the illusion of explanatory depth (Rozenbilt & Keil, 2002). The lack of understanding into what is "beneath the hood" of the GenAI, opens the door to risks when integrating AI and GenAI technologies into high-risk environments, such as the battlefield, emergency response, and disaster areas, among others.

These types of operations are performed in complex and continuously changing environments where every decision, no matter how small, can be a life and death decision. Research into AI at the tactical edge is proving to be a viable way to rapidly evaluate high risk environments, especially when partnered with GenAI technologies for more informed decision-making. These technologies can extend the chances of detecting and identifying threats, obstacles, moving targets, or even missing people, but any tactical edge solution must be sufficiently portable, survivable, and supportable in austere field conditions. Technical solutions must be light and rugged, preferably demanding little power and able to integrate into existing operation equipment (e.g., Nett Warrior, US General Services Administration, 2026; US Army SBIR/STTR Program, 2024). Conforming to the above requirements while optimizing a solution's performance to assist human team members in high-risk situations raises serious challenges. To preserve power and operate on small human-carried devices or small remotely deployed devices, the solution must be efficient and lightweight, and work in a distributed fashion. Solutions with larger AI models and integrating with GenAI technologies require heavier processing that can be difficult to carry or deploy in the field. Therefore, research is underway to identify best practices that might allow relatively high bandwidth reach back capabilities to more robust compute platforms away from the front lines to provide a novel way to integrate with GenAI technologies (Richarson, 2025). Finally, utilizing the modest sensors likely available on an edge compute device may limit the device's detection threshold for any one modality, speaking to a need for leveraging multiple modalities (Zaheer et al., 2023).

Given these technical requirements and limitations, it is even more important that team members, especially those making decisions based on AI-enabled analytics composition, understand the information coming from AI at the tactical edge and the connections to GenAI at a centralized location. Blanket trust in these technologies can open the door to risks from data sources, training, data poisoning, or adversarial attack (Raghavan & Schneier, 2025). Mitigating these risks often requires human experts to rapidly detect violations in AI outputs and recognize when AI crosses into questionable territory, but AI verification may require additional conscious effort on the expert's part (Lee et al., 2025).

Effective human and AI-agentic teaming that results in adequate, appropriate, and trusted operations builds on three major theories coming from human factors: situation awareness (SA), transparency, and situation-awareness based agent transparency (SAT). SA is the *perception* of the environment, *understanding* the meaning behind the perception, and the

ability to *project* future states, allowing for informed decision-making (Endsley, 1995; Endsley & Jones, 2024). Transparency is a design principle ensuring the system's goals, actions, intentions, and limitations are directly observable and understandable by the human team member. When aligned to trust, there are three components, the purpose, process, and performance, of an AI that influence the human team member's trust calibration (Lee & See, 2004). The SAT model (Chen et al., 2018) brings together the foundation of SA and theories of transparency in a three-level approach to quantify AI-agentic design that can be used to mitigate current AI risks and miscalibrations in trust: the agent's current status, actions, and plans (Level 1), reasoning process (Level 2), and agent projection and predictions of future outcome (Level 3). Incorporating the SAT principles builds trust in the GenAI-Edge connection by exposing the way models interact with edge devices.

GENAI AND THE TACTICAL EDGE

The connection between tactical edge data and GenAI is not a new concept. GenAI has shown a great deal of promise in the area of damage assessment and disaster relief following earthquakes (leveraging social media data and image classification), wildfires (edge data coming from optical, thermal, and radar sensors), and even coastline and storm damage following a cyclone (satellite imagery, aerial surveys, and ground observation data; Raj et al., 2025). In a review on GenAI in military operations, GenAI has been identified as an added technical capability to support strategic planning, decision-making, and operational efficiency due to its ability to process and interpret large datasets in real time (Vasankari & Koski, 2025). Both literature reviews highlight the importances of understanding the ethical considerations, bias mitigation, and adversarial vulnerabilities before actively deploying these technologies. That said, these technologies can enable processing at the tactical edge. For example, Swarm-Enabled Hierarchical Intelligent Edge Defense (SHIELD) is a system designed to enable use of convolutional neural networks (CNNs) and recurrent neural networks (RNNs) at the edge, integrated with GenAI and federated learning on the cloud whereby updates from the GenAI can be pushed to the tactical drones as needed (Sarker & Krishnamachari, 2026). In addition, GenAI has been shown to be a viable technology for generating diverse and adversarial simulation scenarios for agents on the tactical edge to learn robust communication and coordination strategies under a wide variety of conditions without risking mission assets, thus improving resiliency at the edge (Ray, 2025).

History and Key Technical Breakthroughs

While a number of resources are available that provide a detailed look into the history of GenAI, this section provides a brief review of key technical advancements that would not have been possible without the access to *big data*, reduction in the requirements for labeling data (i.e., pre-training on unlabeled data and finetuning on specific data), and multimodal data processing. Most GenAI uses some combination of these advancements.

- **RNNs & Long Short-Term Memory (LSTMs; 1980s-90s):** Attempts at modeling sequential data, laying the groundwork for text generation.
- **Generative Adversarial Networks (GANs; 2014):** The GANS framework uses competing generator and discriminator neural networks to improve realistic image generation (Goodfellow et al., 2025).
- **Transformer Architecture (2017):** Transformers use attention mechanisms to process whole sentences at once, allowing for better contextual understanding and scalability (Vaswani et al., 2017).
- **Large Language Models (LLMs) and ChatGPT Release (2022):** The shift toward training on vast, diverse datasets allowed models to perform a wide range of tasks beyond their original training. The public release of ChatGPT demonstrated the capability of GenAI to produce human-like, context-aware, conversational responses, driving mainstream adoption.
- **Reasoning Models (2024–2025):** The introduction of models that utilize chain-of-thought techniques to improve problem-solving

Risks

While these technologies have drastically changed how we process and understand data, there are critical risks that can affect SA, transparency and trust. The first is in the data itself. The performance of an LLM or GenAI is dependent on its underlying data, including quality and size. In many cases of *big data*, there is little information pertaining to how a dataset was created or if it is missing critical metadata (Penedo et al., 2024). In addition, *AI bias* becomes a critical concern. Bias can occur in the data sets themselves (e.g., including direct stereotypical language or imagery in data sets, or sampling bias), bias in design and development through technical constraints (e.g., human labeled data impacted by human cognitive biases, overfitting, and class imbalance), and bias in use (e.g., decisions made; Bode & Bhila 2024). Data can also be purposefully attacked to affect use.

- **Data Poisoning**: Data poisoning happens when an adversary intentionally corrupts the data used to train an AI model by altering training images, misclassification of one vehicle as a different vehicle, or even to ignore the target of interest (Hartle, Mancini, & Kerry, 2025).
- **Adversarial Attacks:** Adversarial attacks exploit blind spots in AI models using intentionally crafted inputs. This differs from data poisoning in that adversarial attacks occur during operations rather than during training (Collins, 2026). For instance, something as simple as placing a sticker on a vehicle or sign causes the AI to fail to recognize it or mislabel that target.

- **Hallucinations:** GenAI models are prone to generate unintended or plainly false responses, a phenomenon known as "hallucination" (Ji et al., 2023). This is used in deception to trigger a waste of resources on a non-existent threat, or to artificially trigger a non-viable action or response.
- **Model Stealing:** This occurs when an adversary probes an AI model to recreate its functionality or steal the model itself. This would allow them to replicate a model, then analyze its weaknesses, develop countermeasures, or put it to their own use, and erase the tactical advantage that our developed AI could provide (Grosse et al., 2024).

USE CASE

This Use Case walks through robust AI mechanisms necessary to take raw tactical edge sensing data and transform it into actionable mission inference.

Step 1: Understanding the User Case

Accurate detection of moving objects, such as people, vehicles, or even debris, can provide critical SA of critical operations whether they are in combat, emergency response, or disaster recovery. No matter the situation, the underlying framework through which intellectuals often articulate desired team behavior is most prominently the Observe, Orient, Decide, and Act (OODA) Cycle (Daniels, 2021). AI provides opportunities to speed up and extend the OODA cycle, enhancing decision making and ensuring increased response and safety (Raska, 2024).

While a large portion of scientific AI research focuses on visual detection and classification, it comes with multiple unique challenges. Terrain, foliage, and adversary obfuscation or camouflage efforts may block visual identification. For example, Carvelli (2025) highlights how success or failure to use natural and artificial visual obfuscation in the Ukraine War has respectively led to significant force preservation and terrible loss of life. As such, non-visual modalities can be used in place of or alongside these visual edge technologies. AI technology can extend a team member's chances of detecting and identifying threats or targets.

One line of effort we define in our example use case identifies less-computationally intensive options whereby alternative sensing could provide adequate detection. In this context, a data science approach can be leveraged to speed up the OODA cycle by creating a pair of multi-modal classification AI models using edge acoustic and seismic data collected from modest sensors. When paired with an inference engine, the composite AI-enabled solution (henceforth referred to as the solution) will quickly and accurately identify targeted vehicles in near-real time. Figure 1 provides a concept for how a human team member or operator might use a trained acoustic and seismic solution to perform initial detection of a vehicle and then cue a mobile AI-enabled asset to monitor a target.

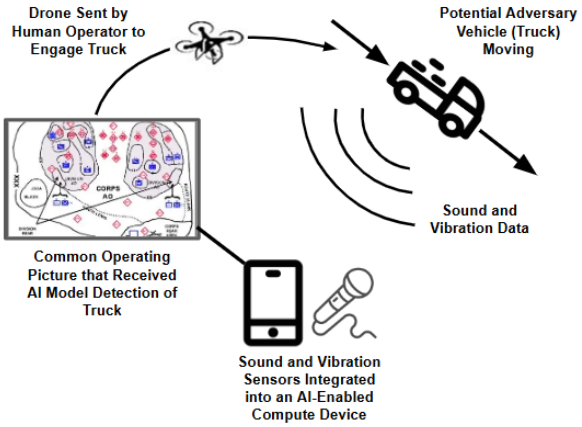

Figure 1: Concept for how a human team member would use a trained acoustic and seismic solution to detect and engage a moving vehicle.

Step 2: Developing a Data Architecture

Setting up a data architecture enables successful integration of AI from the data source all the way through data outcomes. This is a critical step to understanding ground truth and digital signal processing filters to remove noise in data before training can begin. Figure 2 shows a sample data architecture for tactical edge building on the Modern Data Architecture approach (Bornstein, Casado, & Li, 2020).

Several performant private sector and academic acoustic and seismic detection AI algorithms and models exist (Hashima et al. 2025; Liu et al., 2023; Zaheer et al., 2023). These offer a starting point to applying both acoustic and seismic modalities together to minimize the weaknesses of the modest sensors available in edge or embedded devices (Hashima et al., 2025; Liu et al., 2023). Focusing on a lightweight solution enables employment on both simulated as well as real low-compute devices such as mobile phones or a Raspberry Pi.

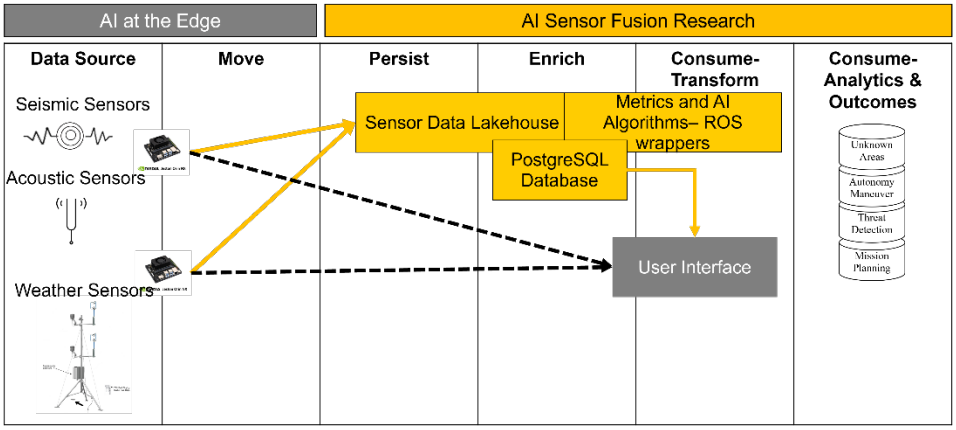

Figure 2: Conceptual data architecture for connecting edge data sources.

Limiting the solution's computational requirements further restrains any power or cooling requirements, essential for edge employment in divergent climate and support environment. While many potential solutions exist, this work uses a multi-pod approach to work through the above developed data architecture. The solution consists of an ingestion pod, signal processing pod, inference pods, and output generation pod supporting two AI models (Figure 3).

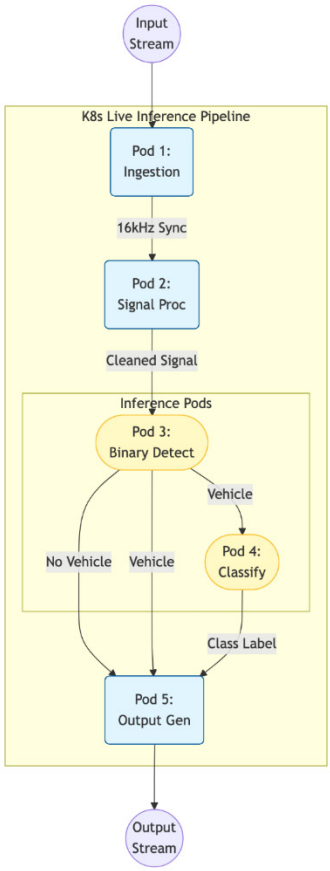

Figure 3: Edge solution diagram for developing effective AI.

The Ingestion Pod is capable of synchronizing two distinct data streams from the Data Source: high-fidelity acoustic data and multi-channel seismic data (comprising a geophone reading and 3-axis accelerometer readings). The acoustic data occurred at either 16 kHz or 1.6 kHz, seismic data at either 200 Hz or 100 Hz, all data was upscaled to 16 kHz before processing (MOD Dataset (Liu et al., 2023) and M3N-VC Dataset; Li et al., 2025). This data is passed to the Processing Pod for noise reduction. The refined signal is first analyzed by a Detection Pod running a binary model; only if a vehicle is detected does the pipeline trigger the Classification Pod identifying the likely vehicle type detected. Finally, the results are formatted by an Output module into Robotics Operating System 2 (ROS2) topics feeding into the Analytics & Outcomes associated with threat detection of the vehicle, autonomy maneuver for the unmanned aerial system for confirmation of the threat, and

broader mission planning. AI models are trained off-solution from historical data pulled from the database offline with tuned models being ported into the inference pods during solution development and updates. The solution team takes special care to efficiently optimize trained model for resource constraints.

Step 3: Integrating AI for Model Training

This step implements the supervised learning models and the iterative loop of training candidate algorithms, validating them against the test set, and freezing the weights of the best performer for deployment. This design follows a conditional, multi-stage pipeline structure tailored for complex, multi-modal signal processing. Initial candidate algorithms for AI model training include CNN (TensorFlow, 2024a), LSTM (TensorFlow, 2024b), and minimalized randomized convolutional kernel transformer (miniROCKET; Demster, Schmidt, & Webb, 2024). These algorithms were chosen for their ability to quickly process time-series data with a high degree of accuracy and relatively low compute requirements.

Step 4: Deployment

This step maps the software logic to hardware, deploying the Docker containers to the Kubernetes cluster and establishing the ROS 2 communication bridge. The solution scales to meet the compute platform and data volume required. Initial testing on a small server indicates that the core pods and AI models could be supported on a device as small as a phone or Raspberry Pi. This offers the prospect of multiple solution instances, independently working in field conditions. With sufficient network connectivity and system integration, these solution instances can provide concurrent queues or indications of activity to a common operating picture. As previously depicted in Figure 1, the SA gained through such a message or symbol (e.g., MIL-STD-2525; Department of Defense, 2014) to a common operating picture can communicate targeting information vital for target engagement using platforms such as unmanned aerial systems.

CONCLUSION

Successful use of GenAI in high-risk environments such as military operations or disaster relief are not limited to the model itself. The data science pipeline with a trust-centric integration approach is key. As shown in the steps of the target detection Use Case above, tactical edge data can actively be used to enable AI for detection and classification. Outputs from the proposed solution even enable mission and path planning for follow on effects from platforms, such as unmanned systems. However, GenAI really comes into play when a human team member needs more information to understand received data or underlying reasoning behind the mission plan, especially when the expert sees a potential risk in the AI output. The above Use Case was a relatively simple example case. When additional data sources come into play to inform

these outcomes, this process becomes even more complex, further impacting human team members trust calibration in the AI. The following provide some guidelines and considerations when integrating GenAI technologies.

- **Asking Questions:** LLMs can be used for the human user to ask questions to better understand the data. To enable this opportunity, a data fabric built early in Step 2 can enable these to function appropriately, especially as new data sources or synthetic data sources are added to the full system.
- **Understanding the Mission:** GenAI for image generation can provide higher fidelity understanding of the mission planning and operations, especially with threat classification and models. However, this can be impacted by data poisoning at both the training and real-time processing stages, leading to a misuse of the system.
- **Code Development:** GenAI for code development can enable rapid data integration and evaluation. While major advancements have been made in this area, review and testing should be conducted for each novel use case.
- **SA:** GenAI embedded into SA platforms can support mission planning by generating and evaluating multiple courses of action based on operational constraints, resource availability, and adversary behavior. These systems propose alternative mission plans, highlight potential risks, and explain the trade-offs between operational strategies. This capability enables humans to explore a wider decision space while retaining final authority over mission execution.
- **Simulation and Synthetic Data:** During training, GenAI can support readiness by producing synthetic scenarios, simulated adversary tactics, and dynamic mission environments for operator rehearsal. Synthetic training environments allow teams to practice responding to complex or rare operational situations that may not frequently occur in real-world exercises. GenAI can help improve familiarity with AI-supported decision systems and strengthen trust calibration through repeated interaction.
- **Teaming:** GenAI can serve as a cognitive interface layer between complex systems and human decision makers. By translating technical outputs into natural language explanations, visual summaries, and risk assessments, generative systems can help ensure that operators maintain easier SA of information technology, AI, and, when sufficiently sensored, physical systems. This role is particularly important in high-tempo operations where rapid comprehension and clear communication can impact reaction to failures and preventative actions with mission critical platforms.

ACKNOWLEDGMENT

This material is partially based upon work supported by the U.S. Army Research Laboratory (ARL) W911NF-17-2-0196 and the U.S. Army Research Office and the U.S. Army Futures Command under Contract No. W519TC-23-C-0045. The authors would like to thank Dr. Damon Conover, ARL, for his mentorship and support. The content of the information does not reflect

the position or the policy of the government and no official endorsement should be inferred.

REFERENCES

Bode, I., & Bhila, I. (2024). "The problem of algorithmic bias in AI-based military decision support systems", in: ICRC Humanitarian Law and Policy. https://blogs.icrc.org/law-and-policy/2024/09/03/the-problem-of-algorithmic-bias-in-ai-based-military-decision-support-systems/

Bornstein, M., Casado, M., & Li, J. (2020). Emerging Architectures for Modern Data Infrastructure. Andreessen Horowitz. https://a16z.com/emerging-architectures-for-modern-data-infrastructure/

Carvelli, M.P. (2025). Invest in battlefield obscuration to win during large-scale combat operations. Military Review, pp. 101–107.

Chen, J.Y.C., Lakhmani, S.G., Stowers, K., et al. (2018). Situation Awareness-based Agent Transparency and Human Autonomy Teaming Effectiveness. Theoretical Issues of Ergonomics Science, Volume 19, No. 3, pp. 259–282.

Collins, A.C. (2026). Garbage in, garbage out? How the monster of AI art reflects human fault, bias, and capitalism in contemporary culture. AI & Society, pp. 1–3

Daniels, O. (2021). "Speeding up the OODA loop with AI: A helpful or limiting framework?" in: Joint Air & Space Power Conference, Power Competence Centre. https://www.japcc.org/essays/speeding-up-the-ooda-loop-with-ai/

Dempster, A., Schmidt, D.F., & Webb, G.I. (2024). miniROCKET [Computer software]. GitHub. https://github.com/angus924/minirocket/tree/main

Department of Defense. (2014). MIL-STD-2525D: Joint Military Symbology. Washington, DC.

Endsley, M.R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. Human Factors, Volume 37, No. 1, pp. 32–64.

Endsley, M.R., & Jones, D.G. (2024). Situation Awareness Oriented Design: Review and Future Directions. Int. J. of Human-Computer Interaction, Volume 40, pp. 1487–1504.

Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, A. (1999) Generative Adversarial Networks. arXiv:1406.2661

Grosse, K., Bieringer, L., Besold, T.R., & Alahi, A.M. (2024). Towards more practical threat models in artificial intelligence security. *USENIX Security, Volume 24,* pp. 4891–4908.

Hartle III, F., Mancini, S., & Kerry, E. (2025). Data poisoning 2018–2025: A systematic review of risks, impacts, and mitigation challenges. Issues in Information Systems, Volume 25, No. 4, pp. 433–442.

Hashima, S., Saad, M. H., Ahmad, A.B., et al. (2025). Effective deep learning aided vehicle classification approach using seismic data. Scientific Reports, Volume 15, No. 22624. https://doi.org/10.1038/s41598-025-01684-x

Ji, Z., Lee, N., Frieske, R., Yu, et al. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, Volume 55, No. 12, pp. 1–38.

Krueger, J. & Dunning, D. (1999). Unskilled and Unaware of it: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments. J. Personality and Soc. Psych., Volume 77, No. 6, pp. 1121–1134.

Lee, H-P., Sarkar, A., Tankelevitch, L., et al. (2025). "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers", in: 2025 CHI Conf. on Human Factors in Computing Systems, pp.1–21. ACM; New York.

Lee, J.D., & See, K.A. (2004). Trust in Automation: Designing for Appropriate Reliance. Human Factors, Volume 46, No. 1, pp. 50–80.

Li, J., Chen, Y., Wang, R., Kimura, T., Wang, T., et al. (2023). "RestoreML: Practical Unsupervised Tuning of Deployed Intelligent IoT Systems," in: DCOSS-IoT, pp. 109–117. IEEE

Liu, S., Kimura, T., Liu, D., et al. (2023). Focal: Contrastive learning for multimodal time-series sensing signals in factorized orthogonal latent space. Advances in Neural Information Processing Systems, Volume 36, pp. 47309–47338.

MIT (2023). Explained: Generative AI. https://news.mit.edu/2023/explained-generative-ai-1109

Penedo, G., Kydlíček, H., Lozhkov, A., et al. (2024). The fineweb datasets: Decanting the web for the finest text data at scale. Advances in Neural Information Processing Systems, Volume 37, pp. 30811–30849.

Raghavan, B. & Schneier, B. (2025). Agentic AI's OODA Loop Problem. IEEE Security & Privacy, Volume 23, No. 6, pp. 80–82.

Raj, A., Shetgaonkar, A., Arora, L., et al. (2025). "AI and Genetic AI Transforming Disaster Management: A Survey of Damage Assessment and Response Techniques", in: IEEE COMPSAC, Toronto, ON, Canada, pp. 1834–1840.

Raska, M. (2024). Reshaping air power doctrines: Creating AI-enabled "super-OODA loops." The Air Power Journal. https://theairpowerjournal.com/reshaping-air-power-doctrines-creating-ai-enabled-super-ooda-loops/

Ray, A. (2025). EdgeAgentX-DT: Integrating Digital Twins and Generative AI for Resilient Edge Intelligence in Tactical Networks. arXiv.2507.21196

Richardson, K. (2025). C5ISR Center Enhances 5G Wireless Network Technology. https://www.army.mil/article/286656/c5isr_center_enhances_5g_wireless_network_technology

Rozenbilt, L. & Keil, F. (2002). The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth. Cogn Sci, Volume 26, No. 5, pp. 521–562.

Sarker, T., & Krishnamachari, B. (2026). "SHIELD: Swarm-Enabled Hierarchical Intelligent Edge Defense for Drone Swarms", in: GENZERO Workshop, M. Andreoni & S. Thankkar (eds.), pp. 56–64. Springer, Singapore

TensorFlow. (2024a). tf.keras.layers.LSTM [Computer software] https://www.tensorflow.org/api_docs/python/tf/keras/layers/LSTM

TensorFlow. (2024b). Simple Audio Recognition: Recognizing keywords. https://www.tensorflow.org/tutorials/audio/simple_audio?hl=en

US Army SBIR/STTR Program. (2024). Development of an Unmanned Aerial Systems Passive Detection, Tracking, and Identification System for Ground Vehicles. https://armysbir.army.mil/topics/development-unmanned-aerial-systems-uas-passive-detection-tracking-identification-system-ground-vehicles/

US General Services Administration. (2026). Contract Opportunity: Acoustic Unmanned Aircraft System Detection and Localization for Dismount Request for Information. https://sam.gov/workspace/contract/opp/a893c9f5eabb41e79502462f9d9d63c1/view

Vasankari, L. & Koski, A. (2025). GenAI in the Military: Trends and Opportunities. Scandinavian Journal of Military Studies, Volume 8, No. 1, pp. 416–434.

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, Volume 30.

Zaheer, R., Ahmad, I., Habibi, D., Islan K.Y., & Phung, Q.V. (2023). A Survey on Artificial Intelligence-based Acoustic Source Identification. IEEE Access, Volume 11, pp. 60078–60108.