Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 60–70
https://doi.org/10.54941/ahfe1008241

Trusting the Chain, Not the Black Box: A Provenance-Based Architecture for Human- Machine Teaming in Military Autonomy

Richard Stone[1], Michael Jenkins[2], and Joseph H. Fernandez[1]

[1]Stone Solutions + Research Collective LLC, Phoenix, AZ 85028, USA
[2]Knowmadics, Inc., Herndon, VA, 20171, USA

ABSTRACT

Military adoption of AI-enabled autonomy depends on more than algorithmic performance. Commanders and operators must be able to determine when an autonomous recommendation is trustworthy, when it is stale or insufficiently corroborated, who or what generated the underlying data, what authority governs the recommended action, and how the action can be audited after execution. This paper presents a human-machine teaming trust model for military autonomy based on accountable provenance rather than opaque model confidence alone. The architecture is an accountable autonomy substrate for multi- domain manned-unmanned teaming in contested environments. Signed mission records, provenance-preserving data custody, target-quality decay, corroboration-gated confirmation, and policy-bound human authorization create a shared operational record between humans and autonomous systems. Large sensor, map, and reconstruction products remain off-ledger behind signed cryptographic commitments, so a military user can inspect the chain of evidence behind a recommendation without every platform carrying exquisite sensing, navigation, or compute. The human-factors contribution is a trust- calibration framework organized around five operator questions: Can I verify the source? Was it altered? Is it current enough for the mission? Is the machine acting within authorized policy and command intent? Can I audit what happened afterward? An operator-facing trust card renders the answers. Simulation results serve as a case study: a stress test with 300 air and ground platforms showed a compact signed index governing an off-ledger data cloud roughly 600 times its size while rejecting tampered data, repairing missing fragments, and prioritizing synchronization under bandwidth limits. The approach shifts the operator's burden from trusting a black-box output to verifying an accountable chain of evidence and authority.

Keywords: Human-machine teaming, Trust calibration, Military autonomy, Data provenance, Autonomous systems, Decision support

INTRODUCTION

Autonomous platforms are arriving in military inventories by the hundred and soon by the thousand: small aircraft, ground vehicles, surface craft, and underwater vehicles cheap enough to lose, each of which senses, navigates, and increasingly recommends. The question that matters at the operator's

Received October 4, 2026; Revised November 15, 2026; Accepted November 22, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

console is not how accurate the recognition model was on its test set. It is whether this recommendation, now, should be acted on, and whether that decision can be defended afterward. A model that reports a single scalar confidence gives an operator no way to separate a track observed twelve seconds ago by three independent platforms from a track reported once, by one platform, four minutes ago. Both can carry the same number.

Two structural trends make this worse. The most capable sensing, navigation, and targeting technologies are concentrated on a few exquisite and expensive platforms, while the geography of a peer conflict demands mass, persistence, and survivability that a handful of exquisite platforms cannot supply. A centralized architecture that relays everything through protected rear nodes fails at the denied, partitioned, emission-controlled edge where the decision has to close. Exquisite capability on every platform is too expensive to field and too precious to risk; cheap mass is blind and disconnected the moment it matters.

The architecture summarized here reframes that problem. Capability does not need to be on the platform; it needs to be reachable by the platform, and the data it consumes needs to be trusted because it is signed, not because the network carrying it is secure. The same property answers the human-factors problem: if every observation, corroboration, classification crossing, and authorization is a signed record on a shared, tamper-evident substrate, the operator no longer has to rely solely on the black-box output. The operator can inspect the evidentiary chain behind it. The paper contributes the substrate whose mechanisms produce that evidence, a trust-calibration framework organized around five questions an operator must be able to answer before acting on a recommendation, the trust card that renders those answers, and a simulation case study, including where the mechanisms fail. The claims are kept at their levels of maturity: the framework is conceptual, the mechanisms are engineering, the simulation results are bounded to declared conditions, and the trust card is a design hypothesis not yet evaluated with operators.

WHY CONFIDENCE ALONE MISCALIBRATES TRUST

Trust in automation has to be calibrated rather than maximized: over-trust produces misuse, under-trust produces disuse, and both degrade joint performance (Lee and See, 2004; Parasuraman and Riley, 1997). Automation bias follows when an aid is usually right and its output is easy to accept unchecked (Parasuraman and Manzey, 2010). Learned trust is shaped by what an interface lets an operator verify, and transparency and communicated uncertainty are among the strongest predictors of calibration (Hoff and Bashir, 2015). A scalar score is also an impoverished basis for situation awareness: the operator needs to perceive the elements behind a recommendation, comprehend their status, and project their implications (Endsley, 1995), and a single number carries none of that structure. Reviews of human- autonomy teaming identify the operator's ability to interpret and verify an autonomous teammate's output as a persistent, unresolved problem (Chen and Barnes, 2014; O'Neill et al., 2022), operational trust moves faster than a static display can track (Rindfuss et al., 2025), and test and evaluation of military AI has to treat the human as part of the system (Helmer et al., 2024).

Explainable AI largely pursues trust by exposing model-internal rationale (Miller, 2019). The approach here is complementary: rather than explaining how a model reasoned, it exposes the evidentiary chain behind what the model asserts, independent of the model's internals, which is what the accountability literature requires when it asks that a decision process be verifiable after the fact (Kroll et al., 2017). The W3C PROV model formalizes provenance as a machine-verifiable record of the entities, activities, and agents behind data (Moreau and Groth, 2013); the signed mission record specializes it for perishable, safety-relevant claims. Subjective logic supplies the uncertainty representation (Jøsang, 2016), and partition-tolerant DAG ledgers for robotic swarms (Tran et al., 2019) are treated as foundational. To our knowledge, prior work does not combine these elements into a single operator-centered framework. That combination is the contribution.

AN ACCOUNTABLE AUTONOMY SUBSTRATE

The architecture relocates trust from the individual platform to a shared, signed record (Figure 1). Six mechanisms produce the evidence the framework requires.

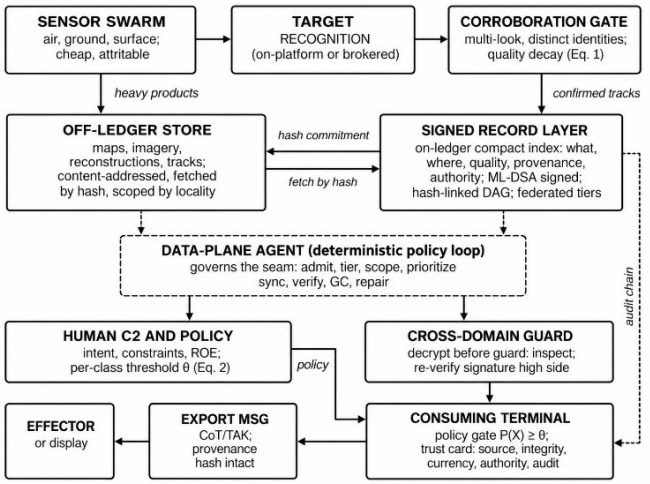


Figure 1: Overall system: platforms produce signed records, heavy data lives off-ledger and is fetched by hash, and the export message carries the provenance hash to the display or effector.

Signed record layer. The foundation is a permissioned, append-only, tamper- evident record shared by vetted nodes. Records are small structured objects, each signed and linked by parent hash into a directed acyclic graph that admits parallel writes, tolerates partition, and reaches Byzantine-fault-tolerant agreement. Access is governed by cryptographic identity, a valid certificate chained to a trusted authority rather than membership of a particular network, and signatures use the post-quantum ML-DSA standard (NIST, 2024). Trust rides with the data and is verified locally, so the substrate works in a disconnected swarm with no central server. A federated hierarchy of mission, theater, and strategic tiers keeps consensus bounded and maps

visibility onto classification and coalition-release boundaries. None of this is claimed as novel.

On/off-ledger seam. A ledger that tried to hold maps, imagery, and full reconstructions would choke at the scale of a real force. The substrate commits only a small signed index stating what exists, where, of what type, at what quality, and with what provenance; the heavy data behind each entry lives in a content- addressed off-ledger store, fetched by hash only when and where a consumer needs it. A deterministic, auditable data-plane agent governs the seam. It is a policy loop, not a generative model in the decision path: it admits and tiers observations, scopes replication by locality, prioritizes scarce bandwidth toward the most time-urgent tracks, verifies that every index entry resolves to a valid stored fragment, garbage- collects, repairs gaps after partition, and rejects any fragment whose content does not match its committed hash. The seam is also what lets a cheap platform borrow capability it does not carry: a camera-only drone localizes against a signed map cloud a mapping platform built, and heavy compute is brokered from the rear as an authenticated service checked on the platform for signature, freshness, and feasibility before use (Fernandez, 2026).

Corroboration-gated confirmation and target-quality decay. Raw detections become confirmed tracks only after multi-look corroboration from distinct signed identities, which rejects false tracks that a single report, however often repeated, could otherwise create. Each confirmed track carries a target- quality score that decays with time, quickly for movers, and is refreshed by a new look; falling below threshold drives automatic re-tasking. The decay runs toward uncertainty rather than falsehood. For supporting and refuting evidence counts r and s, decay constant λ, prior weight W, base rate a, and elapsed time Δt, projected confidence is

$$P(X, \Delta t) = b(\Delta t) + a \cdot u(\Delta t), \text{ with } r(\Delta t) = r \cdot e^{-\lambda \Delta t},\ s(\Delta t) = s \cdot e^{-\lambda \Delta t},\ b(\Delta t) = r(\Delta t) \,/\, (r(\Delta t) + s(\Delta t) + W),\ u(\Delta t) = W \,/\, (r(\Delta t) + s(\Delta t) + W) \quad \text{(Eq. 1)}$$

As Δt grows, uncertainty approaches one and confidence converges to the base rate rather than to zero: a stale track is treated as unknown, not as wrong (Jøsang, 2016). P(X) should not be confused with the opaque model-confidence value criticized above. It is an evidence-derived projected probability computed from explicit supporting and refuting observations whose sources, corroboration, age, and provenance remain independently inspectable. The scalar therefore summarizes one property of the evidence, currency-weighted support, rather than replacing the evidentiary chain on which operator trust is calibrated.

Cross-domain provenance. A record encrypted at the edge is decrypted to inspectable plaintext before it reaches the accredited cross-domain guard; the high side re-verifies the ML-DSA signature, and the provenance chain survives the crossing intact. It is the signature that crosses the guard and carries provenance, so chain of custody from sensor to effector is continuous across classification boundaries.

Coarse-resolution verification. A parent tier must trust and act on a child's heavy off-ledger data without pulling it up the hierarchy, which would defeat the seam and violate need-to-know. The heavy data is transformed

into coarse, mid, and fine frequency bands; the child sends the small coarse band upward and commits, on the ledger, a signed hash of the full coefficient set. The parent verifies the coarse band against that commitment and acts on it without downloading the fine payload, which stays need-to-know at the edge. The limit has to be stated plainly: the commitment proves integrity, meaning the data was not altered after signing, not the truthfulness of the signer's original data; full assurance requires an audit pull or a trusted child.

Policy-bound human authorization. Confidence and authorization are architecturally separated. The substrate computes a confidence value $P(X)$, but a policy-owned threshold θ, set by accountable human policy and not by the model, determines whether a recommendation is surfaced for a given class of action.

Following the minimum-risk decision rule for two actions with correct decisions assigned zero cost (Duda, Hart and Stork, 2001), a recommendation surfaces only if

$$P(X) \geq \theta, \text{ with } \theta = C_{FP} / (C_{FP} + C_{FN}) \qquad \text{(Eq. 2)}$$

where C_{FP} and C_{FN} are policy-defined costs of a false positive and a false negative, so identical confidence values may or may not surface depending on the stakes policy assigns. Authorization is encoded in the record and travels with it. The architecture can support multiple human-control postures, ranging from human- required authorization to policy-bounded autonomous execution, depending on the action class and applicable operational policy. The architecture prescribes no preferred posture. For autonomous or semi-autonomous weapon systems, DoDD 3000.09 independently requires the system to permit commanders and operators to exercise appropriate levels of human judgment over the use of force. The AI layers are bound by the same discipline: they retrieve and present authenticated data, and a generative model may phrase what was retrieved but never add to it. It is a librarian, not an oracle, and it stays out of the use-of-force decision.

A TRUST-CALIBRATION FRAMEWORK: FIVE OPERATOR QUESTIONS

Operator trust in an autonomous recommendation is not a single judgment. It is a composite of distinguishable questions, each of which the substrate answers directly rather than leaving implicit in a confidence score: Can I verify where this came from (source)? Can I determine whether it was altered (integrity)? Can I assess whether it is current enough for the mission (currency)? Can I confirm the machine is acting within authorized policy and command intent (authority)? Can I reconstruct what happened afterward (audit)? Table 1 maps each question to the mechanism that answers it and the evidence the operator sees. One boundary matters: a signature establishes that a record is unmodified since publication, not that its originating sensor was accurate or uncompromised; that concern belongs to source and corroboration.

The framework does not require trusting the system's internal reasoning, only verifying the evidentiary chain behind what it asserts, a narrower and more tractable standard than explaining model internals and the one already

implicit in autonomy policy: requiring appropriate human judgment over the use of force is, in substance, an authorization and auditability requirement, which Q4 and Q5 satisfy by construction.

Table 1: The five operator questions, the mechanism that answers each, and the evidence rendered.

Question	Mechanism	Evidence Rendered
Q1 Source	Signed record; corroboration gate over distinct identities	Publisher identity; corroborating platforms, including cross-domain
Q2 Integrity	Content hash bound to ML-DSA signature; hash-linked chain	Verification status; tamper flags since publication
Q3 Currency	Target-quality decay and refresh (Eq. 1)	Evidence-derived probability; time since last look
Q4 Authority	Policy-owned gate (Eq. 2); authorization in the record	Action class and threshold; inform- only or delegated; geofence and tasking checks
Q5 Audit	Append-only signed provenance chain	Link into the full chain; signed policy object for the gate decision

TRUST CARDS: AN OPERATOR-FACING INTERFACE CONCEPT

The trust card (Figure 2) renders the five questions to an operator. Each card is scoped to a single track or recommendation and closes with a recommended action framed as informing rather than authorizing. The card is generated from the same export message that feeds existing tactical displays, so the provenance hash it shows is the one carried in the record. Field order is deliberate: source and integrity come first, because an unverifiable or altered record makes everything below it moot; currency precedes authority, because staleness can change which threshold applies; and the recommended-action line is visually separated from the diagnostics above it to discourage treating the recommendation as a conclusion. Each of these is a design intention, not a measured effect.

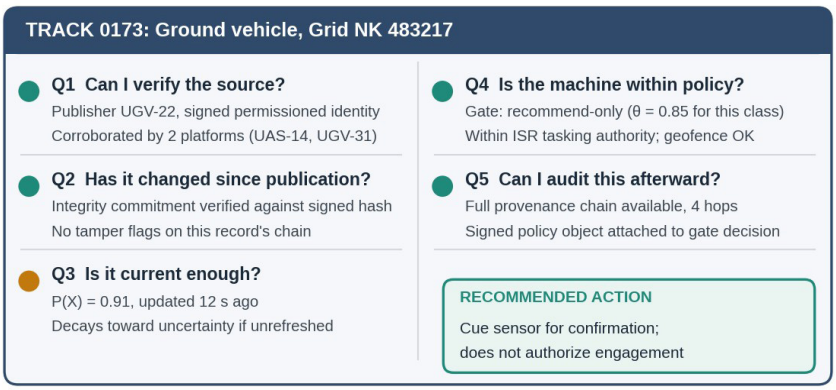


Figure 2: Trust card interface concept, organized around the five operator questions. The recommended action informs; it does not authorize.

The card is a design concept, not a validated interface. Layout, density, terminology, and the status marks have not been evaluated with operators, and the risk that a mostly green card could itself become a heuristic for acceptance, reproducing the automation bias the interface is intended to mitigate. The planned operator study therefore treats appropriate reliance and decision latency as empirical outcomes rather than assumed benefits.

SIMULATION CASE STUDY

Two bodies of simulation evidence are used: architecture-level stress tests from the project record, run on runnable components with real post-quantum signatures where noted, and a kernel-level discrete-event Monte Carlo study that isolates the trust mechanisms under an explicit threat model. Both are existence proofs under declared conditions, showing that the mechanisms behave as designed when pushed, including where they fail, not that the scenarios match a fielded environment.

Architecture-Level Stress Tests

Stress Test A exercised the data-plane agent against 200 drones and 100 ground robots over 259 square kilometers. A 44.8 MB signed on-ledger index governed a 26.7 GB off-ledger cloud, a ratio of 1:596 that improves with scale, which is the architecture's central scaling claim. The consistency seam held under continuous churn, with fragment-store failures repaired, orphans garbage-collected, and every tampered fetch rejected by content addressing. The binding constraint was radio bandwidth rather than ledger or agent compute. One finding has to be reported rather than smoothed over: the synchronization backlog grew without bound under offered load. The agent correctly spent bandwidth on decaying movers while static updates queued, and a production system must age out or shed that backlog. That is a required refinement, not a solved problem.

Coarse-resolution verification was tested with real ML-DSA signatures across five pre-registered hypotheses of 400 trials each. Clean commitments verified coarse-only, coarse-band and localized tampering were caught at the parent without pulling the fine payload, post-signing edits to the fine band were caught on audit because the signed full-coefficient hash binds it, and forgery without the child's private key was rejected, each in 100 percent of trials. A decision- sufficiency sweep found that at 8 by 8 coefficients, one sixty-fourth of the payload, parent tasking decisions matched the full-data decision about 90 percent of the time, while targeting to a precise aimpoint was never coarse-sufficient, which correctly keeps fine detail at the edge. The signed record spine itself is runnable with real post-quantum signatures; security under a real adversary and consensus under adversarial networks remain asserted rather than simulated.

Kernel-Level Monte Carlo Study

To isolate the trust mechanisms, a discrete-event Monte Carlo simulation modeled 300 unmanned platforms (100 aerial, 100 ground, 100 maritime), 60 tracked entities, 200 one-second timesteps, and 30 seeded trials per

experiment, every statistic reproducible from the code and configuration. Six declared threats were modeled: coordinated Sybil clusters synchronizing phantom reports to mimic multi-source corroboration, opportunistic spoofers, stale-claim exploitation, tampering and replica loss in the payload tier, and communications denial as periodic contention windows. Signature forgery, key compromise, and physics- level sensor deception were out of scope.

A track was confirmed under the gated policy only once decayed evidence from distinct signed identities held $P(X)$ at or above $\theta = 0.70$ for 15 consecutive timesteps; a naive baseline confirmed on any single report. The adversary fielded five coordinated clusters of three to six identities plus fifteen opportunistic spoofers, about 12.8 percent of the fleet. In the payload tier, 240 payloads in eight fragments with three-way replication suffered continuous tampering and loss while coarse-first hash checks against the signed index (unit cost, with targeted repair) competed against full-payload retrieval (40 times the cost, no repair path) for one budget. In the synchronization layer, record value was made perishable and a value-priority scheduler was compared with first-in-first-out (FIFO) under the same contention budget. Table 2 reports the headline results; Figure 3 shows one track through a 115-timestep observation gap.

Table 2: Kernel-level results, mean ± SD across 30 seeded trials.

Metric	Trust Mechanism	Comparison Baseline
True track confirmed	92.4 ± 3.2 percent	100.0 percent
Phantom track falsely confirmed	35.8 ± 8.9 percent	100.0 percent
Confirmation latency (timesteps)	20.35 ± 2.32	1.29 ± 0.10
Faulty replicas detected (count)	9,129 ± 97	1,256 ± 26
Bad replicas at end of run	4.57 ± 0.37 percent	81.98 ± 0.49 percent (no repair)
Total perishable value delivered	14,412 ± 273	13,612 ± 249 (FIFO)
Critical items backlogged at window close	0.00	71.02 ± 3.19 (FIFO)

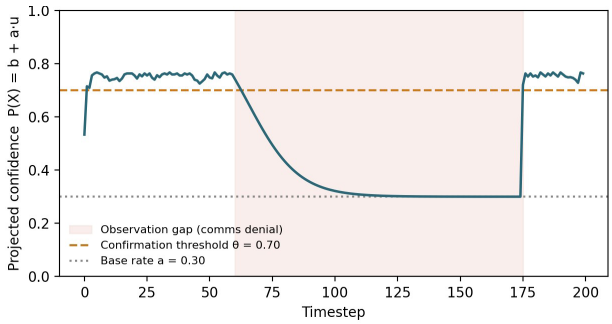


Figure 3: Target-quality decay and refresh for one track across a 115-timestep observation gap: confidence decays toward the base rate a = 0.30, crosses below $\theta = 0.70$ so the stale track de-confirms, and recovers on refresh. Illustrative trace, not an aggregate statistic.

Gating cut false confirmation by a factor of 2.8 under the blended adversary at two honest costs: latency dominated by the 15-step dwell, and a 7.6 percent miss rate on sparsely observed entities whose evidence aged in transit. Coarse hash- checking detected 7.3 times as many faulty replicas as full-payload retrieval under the same budget, and index-guided repair held end-of-run corruption to 4.57 percent against 81.98 percent unrepaired. Value-priority scheduling delivered 5.9 percent more value than FIFO and closed every contention window with an empty critical backlog. Sensitivity sweeps located the failure boundaries: purely coordinated adversaries succeeded 84 percent of the time and clusters of four or more identities defeated identity counting completely, a structural limit only an identity-attestation layer can close; the repair budget, run near full utilization, collapsed to 32 percent corruption at 1.5 times the baseline fault rate; and value- priority protected critical traffic only while bandwidth exceeded critical demand.

DISCUSSION

Read against the claim levels set out in the introduction, the results support the engineering and bounded-simulation claims: the mechanisms hold under declared stress, at measured costs in latency and misses, and their failure boundaries are located. The framework is conceptual, motivated by this evidence rather than proven by it, and the trust card remains a design hypothesis with no evidence yet on operator trust, workload, or decision speed.

The design pattern does not remove the operator's judgment; it changes what the judgment is about. Instead of deciding whether to believe a number, the operator reads a chain: who saw it, who corroborated it, whether it was altered, how old it is, what policy allows, and whether it can be reconstructed later. Three consequences follow. Workload can rise, because five checks are more than one; the card front-loads the two that most often end the inquiry so an operator can stop reading early when a record fails. Failure modes become legible: a track that de- confirms because its evidence aged is visibly different from one that was never corroborated. And machine-speed functions run inside auditable policy bounds while mission intent and use-of-force decisions stay with humans. Machine-speed functions can therefore operate within auditable policy bounds while mission intent and use-of-force judgments remain attributable to accountable humans. These properties support, but do not establish, the human-judgment and auditability requirements relevant to law-of-war review.

Limitations

The simulations are parameter-driven: observation models, adversary behavior, fault rates, traffic, and transport latencies are declared rather than measured, and the dwell and threshold values are policy parameters against which the miss rate trades. Two human gates remain that no simulation closes: a distributed-systems engineer must pressure-test consensus under real adversarial, partitioned networks, and a cryptographer must review the edge-encryption design, especially revocation under capture. The trust

card has not been evaluated with operators; the planned study compares it within-subjects against a scalar-confidence display, measuring appropriate reliance and decision latency. Delegated operational authority is out of scope, and Directive 3000.09 governs weapon systems specifically; this paper claims no compliance with any certification pathway.

CONCLUSION

Calibrated operator trust is a critical human-factors requirement for effective military human-autonomy teaming, and a scalar confidence value provides an incomplete basis for that trust. This paper has presented a trust model for human-machine teaming organized around five operator questions, source, integrity, currency, authority, and audit, each answered directly by an accountable- provenance architecture rather than inferred from a black-box output. Cheap platforms consume exquisite capability across a signed seam while every observation, corroboration, classification crossing, and authorization stays attributable end to end. Simulation results show the mechanisms holding under declared stress and where they break, and the trust card renders the framework in operator terms as a hypothesis for human-subjects evaluation. Humans retain judgment over mission intent and use of force; autonomous systems execute sensing, synchronization, re-tasking, and data management at machine speed within auditable policy bounds. Future work includes the operator study, an identity-attestation layer to close the coordinated-spoofing gap, and validation against operational rather than simulated data.

ACKNOWLEDGMENT

Generative AI disclosure: Claude (Anthropic) assisted with consolidating the authors' architecture documents, white papers, and simulation records into the conference format and with reviewing the draft for internal consistency. All architecture, study design decisions, simulation results, and technical content are the authors' own, and the authors take responsibility for the final text.

REFERENCES

Chen, J.Y.C. and Barnes, M.J. (2014). Human-agent teaming for multirobot control: A review of human factors issues. IEEE Transactions on Human-Machine Systems, 44(1), pp. 13–29.

Department of Defense (2023). Directive 3000.09: Autonomy in weapon systems. Washington, DC: Office of the Under Secretary of Defense for Policy.

Duda, R.O., Hart, P.E. and Stork, D.G. (2001). Pattern classification. 2nd ed. New York: Wiley.

Endsley, M.R. (1995). Toward a theory of situation awareness in dynamic systems. Human Factors, 37(1), pp. 32–64.

Fernandez, J.H. (2026). Trusted capability brokerage for autonomous systems: Capability without carriage. Unpublished concept white paper.

Helmer, D., Boardman, M., Conroy, S.K., Hepworth, A.J. and Harjani, M. (2024). Human-centred test and evaluation of military AI. arXiv preprint arXiv:2412.01978.

Hoff, K.A. and Bashir, M. (2015). Trust in automation: integrating empirical evidence on factors that influence trust. Human Factors, 57(3), pp. 407–434.

Jøsang, A. (2016). Subjective logic: A formalism for reasoning under uncertainty. Cham: Springer.

Kroll, J.A., Huey, J., Barocas, S., Felten, E.W., Reidenberg, J.R., Robinson, D.G. and Yu, H. (2017). Accountable algorithms. University of Pennsylvania Law Review, 165(3), pp. 633–705.

Lee, J.D. and See, K.A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), pp. 50–80.

Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence, 267, pp. 1–38.

Moreau, L. and Groth, P. (2013). Provenance: An introduction to PROV. Synthesis Lectures on the Semantic Web: Theory and Technology, 3(4), pp. 1–129.

National Institute of Standards and Technology (2024). Module-lattice-based digital signature standard. FIPS 204. Gaithersburg, MD: NIST.

O'Neill, T., McNeese, N., Barron, A. and Schelble, B. (2022). Human-autonomy teaming: A review and analysis of the empirical literature. Human Factors, 64(5), pp. 904–938.

Parasuraman, R. and Manzey, D.H. (2010). Complacency and bias in human use of automation: an attentional integration. Human Factors, 52(3), pp. 381–410.

Parasuraman, R. and Riley, V. (1997). Humans and automation: use, misuse, disuse, abuse. Human Factors, 39(2), pp. 230–253.

Rindfuss, A., Leary, S., Dutta, P., Chen, R., Clark, T.K., Kong, Z. and Hayman, A.P.A. (2025). Modeling trust and its dynamics from physiological signals and embedded measures for operational human-autonomy teaming. Frontiers in Robotics and AI, 12, 1624777.

Tran, J.A., Ramachandran, G.S., Shah, P.M., Danilov, C.B., Santiago, R.A. and Krishnamachari, B. (2019). SwarmDAG: A partition tolerant distributed ledger protocol for swarm robotics. Ledger, 4(S1), pp. 25–31.