12th INTERNATIONAL SCIENTIFIC CONFERENCE ON DEFENSIVE TECHNOLOGIES
OTEH 2026
Serbia, Tara, 21 - 24 October 2026

AI-ASSISTED DECISION SUPPORT IN COMMAND-AND-CONTROL SYSTEMS: APPLYING THE SENSE–UNDERSTAND–DECIDE–ACT (SUDA) CYCLE

IVAN VULIĆ, MILJAN VUČETIĆ, ILIJA POPADIĆ, RANKO PETROVIĆ, ZORAN BANJAC
Vlatacom Institute of High Technologies, Milutina Milankovića 5, Belgrade, Serbia

Abstract: *Modern command-and-control (C2) systems face a paradox: the volume of sensor data grows exponentially while the time available for decision-making shrinks. This paper presents an architectural framework for AI-assisted decision support in C2 systems based on the Sense–Understand–Decide–Act (SUDA) cycle and demonstrates it on a deployed system: the Vlatacom C7ISTAR solution, in which the ORAO command platform carries the ARES (Algorithmic Reconnaissance and Execution System) three-tier AI core, in operational service protecting national critical infrastructure since 2024. Each phase is mapped to a technology layer and characterized by measured quantities: a panoramic thermal sensor covering the full circle in 2.3 s; an edge-AI stage above 100 TOPS with sub-100 ms inference latency, reducing about fifty raw detections per second to approximately one plausible target; and an agile electro-optical turret classifying a cued object in under two seconds. We argue that decision advantage arises not from replacing the human decision-maker but from restructuring the C2 information flow around the SUDA cycle, so that the path from sensor to decision-maker, and from decision to effect, becomes short, traceable and measurable.*

Keywords: *command and control, decision support, artificial intelligence, SUDA cycle, counter-UAS.*

1. INTRODUCTION

Sensors have proliferated across every domain of the modern battlefield—radar, electro-optical and infrared imaging, radio-frequency and acoustic sensing, unmanned platforms, and networked reporting—so a contemporary command post receives orders of magnitude more data than any human staff can assimilate in the time available. At the same time the velocity of the threat has increased: unmanned systems in large numbers, loitering munitions, electronic attack and cyber effects compress the interval between first detection and required response from hours to minutes, and in terminal defense to seconds.

Hence the central paradox of this paper: commanders have never had more information, and never less time in which to use it. The side that resolves it obtains decision advantage — the ability to close its own decision cycle faster, and with better judgment, than the adversary. The classical articulation is Boyd's observe–orient–decide–act (OODA) loop [1], refined for C2 analysis by Brehmer [2]; Endsley's theory of situation awareness [3] captures the quality dimension. Command-support terminology has evolved in step: from C2 through C4ISR — within which geographic information systems first brought spatial decision support to the operational picture [4] — to the extended C5ISR, C6ISTAR and C7ISTAR designations that successively admit cyber defense, combat systems, and collaboration and interoperability [5]; C7ISTAR, the fullest of these, is adopted here as the designation of the unified sovereign command-and-mission portfolio that the described solution implements.

Artificial intelligence can now materially contribute to both the speed and the quality of the cycle, yet its integration into C2 raises questions as much doctrinal as technical — how much authority may be delegated to a machine, how a recommendation is explained to the operator who must act on it, how models operating on classified data are kept under sovereign control — and recent scholarship agrees that acceleration is valuable only while human judgment remains authoritative within the loop [6]. Beneath these lies the practical question this paper takes as its organizing principle: how long is the path from sensor to decision-maker, what happens to the information along the way, and how much of the journey can be entrusted to a machine.

The paper defines the SUDA cycle and positions it against the OODA loop, situation-awareness theory and the JDL data-fusion model (Section 2); maps each phase onto a technology layer (Section 3); instantiates the framework on the deployed Vlatacom C7ISTAR solution, with measured quantities per phase (Section 4); analyses human–machine teaming and the constraints of edge inference and sovereignty (Sections 5 and 6); and traces a complete engagement with its cycle-time budget (Section 7). Section 8 concludes.

2. THE SUDA CYCLE AS A C2 FRAMEWORK

The SUDA cycle structures a C2 system's information flow into four continuously repeating phases. In Sense, data are acquired from all available sources. In Understand, they are fused into a single coherent representation — the common operational picture (COP) — and enriched with machine-derived interpretation: classification, anomaly detection, and short-term prediction. In Decide, the system generates feasible courses of action (COA), evaluates them against constraints, and presents them to the human decision-

Corresponding author: I. Vulić;
E-mail: ivan.vulic@vlatacom.com

DOI: 10.5937/OTEH26102V

maker. In Act, the selected course is translated into tasking of effectors and sensors, execution is monitored, and observed effects feed the next iteration; Figure 1 summarises the cycle.

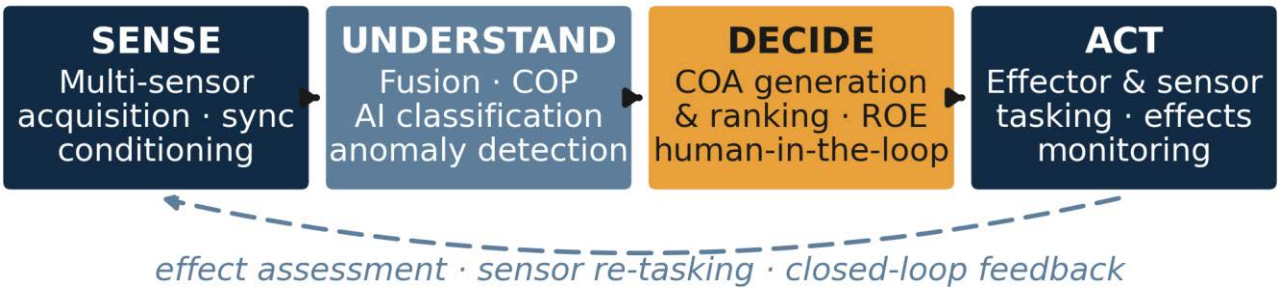


Figure 1. The SUDA decision cycle with the technology layer associated with each phase

SUDA deliberately restates Boyd's loop in system-engineering form: Observe maps to Sense; Orient — the phase Boyd considered decisive — to Understand, where fusion and machine learning are concentrated; Decide and Act retain their names. The gain is mapping discipline: each phase becomes an identifiable set of services with interfaces and latency budgets, and cycle time — from a change in the real world to a completed response — becomes a key performance indicator, decomposable into per-phase terms that Section 4 measures on a deployed system.

The framework aligns with established theory and doctrine. Endsley's three levels of situation awareness [3] span Sense and Understand; the JDL data-fusion model [7] structures the Understand phase—historically the bottleneck of the C2 cycle and therefore where AI offers the greatest compression —the central design thesis of this paper. The Joint All-Domain Command and Control (JADC2) concept likewise organizes capability around sense, make sense, and act — matching Sense, Understand, and the combined Decide–Act — and requires machine analysis to be explainable, confidence-tagged, and accessible only to authorized users [8], [5]. SUDA is thus not a competitor to these formulations but their engineering expression.

3. MAPPING SUDA PHASES TO TECHNOLOGY LAYERS

Figure 2 presents the functional architecture resulting from mapping the SUDA phases onto technology layers: read upwards from the base, it is the sensor-to-decision-maker path drawn as a stack, with the cryptographic rail protecting the observation at every station.

3.1. Sense — multi-sensor acquisition

The Sense layer comprises the heterogeneous sensor set — radar, electro-optical and infrared imaging, radio-frequency and acoustic sensors, unattended ground sensors — together with structured reporting from deployed units and, increasingly, open-source (OSINT) feeds. Three engineering requirements dominate. First, time synchronization and geo-registration: fusion is only as good as the spatial and temporal alignment of its inputs, so every observation must carry a trusted timestamp and position reference. Second, standardized metadata: sensor identity, mode, confidence, and provenance must accompany every detection so the evidence can be weighed correctly; that trust begins with authenticating the sensors themselves, for which the authors' earlier Internet-of-Battlefield-Things model provides a basis [9]. Third, data conditioning at the edge, so that constrained links carry information rather than raw redundancy. The layer must also degrade gracefully: losing individual sensors or links should reduce coverage, not collapse the picture.

3.2. Understand — fusion, the COP and AI services

The Understand layer transforms conditioned observations into shared meaning. Its foundation is multi-sensor fusion in the sense of the JDL model [7]: associating detections into tracks, correlating tracks across sensors, and maintaining a single consistent COP that all echelons consume. Fusion may occur at the detection level or the track level; practical systems combine both, trading communication load against fusion gain [10], and at the semantic end of this range the Institute has applied fuzzy-relational methods to fuse heterogeneous surveillance data [11]. The integrity of the COP's geospatial foundation must be protected in its own right [12].

On top of the fused picture, AI services operate: classification models assign object types; anomaly-detection models learn an area's pattern of life and flag deviations that escape threshold-based rules; predictive models project track behavior, supporting projection in Endsley's sense [3]. Two properties are essential: every machine-derived statement carries explicit uncertainty — confidence and provenance, surfaced to the user — and the services are modular and retrainable, components with a lifecycle.

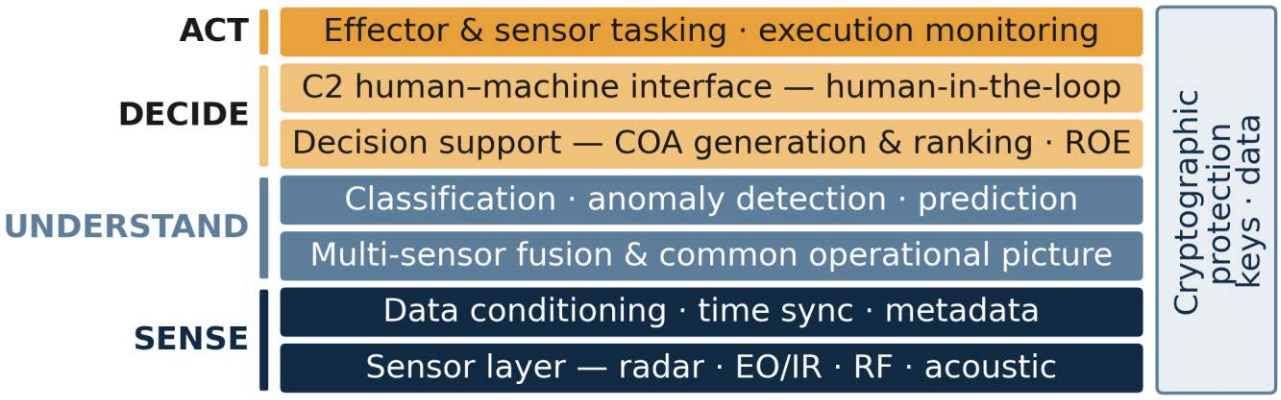


Figure 2. Functional architecture of AI-assisted C2, organized by SUDA phase, with cryptographic protection as a cross-cutting layer

3.3. Decide — decision support with the human in the loop

The Decide layer converts understanding into options: threat evaluation, COA generation, and ranking against explicit criteria — effectiveness, resource consumption, risk, deconfliction and, critically, compliance with the rules of engagement (ROE), encoded as machine-readable constraints so that no proposed option violates them by construction; what-if simulation lets the operator probe an option before committing. Large language models have recently been demonstrated as generators of candidate courses of action at machine speed [13]; the same literature documents their unreliability under adversarial and escalatory conditions [14] — which is precisely why this layer is designed so that the machine proposes and the human disposes. Ranked options are presented with their rationale — factors, contributing sensors, confidence — and every recommendation and human selection is logged, creating the audit trail on which accountability and institutional learning depend.

3.4. Act — tasking, monitoring and feedback

The Act layer translates the selected course of action into tasking of effectors and, equally, of sensors: closing the loop includes directing collection where confirmation of effect is needed. Execution is monitored against the plan, and effect assessment — whether the action achieved the intended result — is fed back into Sense and Understand, updating the picture and, over time, the models themselves. In a realistic system, many SUDA cycles run concurrently, from seconds in point defense to hours at the operational level.

4. REALISATION: THE ORAO C7ISTAR PLATFORM

The framework of Section 3 is realized in ORAO, the command-and-control platform of the Vlatacom Institute (the name is proper to this software platform, unrelated to the aircraft): a scalable micro-services system that integrates remote sensing equipment, maintains a common operational picture and lets the operator act in the field remotely. ORAO has been in operational service since 2024, protecting a national critical energy-infrastructure installation against small commercial and FPV drones — the class current counter-UAS surveys identify as hardest to counter at such sites [15], [16]. It rests on a decade of the Institute's research across the SUDA layers: electro-optical platforms and video processing [17], [18], object tracking and AI-based detection [19], [20], [21], fusion of heterogeneous sensor data [11], and cryptographic protection of operational data [22], [23].

ARES, the Algorithmic Reconnaissance and Execution System, supplies the platform's intelligence— a modular AI core conceived as a software brain over the existing sensor base, organized in three tiers that map almost one-to-one onto the SUDA phases. ARES Edge infers on a processing module at the sensor itself, classifying targets in under a second without transmitting a high-bandwidth video stream — the transition from Sense to early Understand, viable on constrained or intermittent networks. ARES Core is the operational tier: it fuses encrypted telemetry of the order of ten thousand events per second from several sites into a single COP and applies a unified ontology that recognizes behavioral patterns rather than isolated events, carrying Understand and preparing Decide. ARES Strategic is the command tier, coupling the picture to large-language-model agents through which a commander obtains reports and candidate courses of action by voice or text — the pattern demonstrated by [13], here on an operational picture rather than a simulation, under the human authority and audit trail of Section 3.3. Across all tiers, AI is combined with rule-based events and actions in a supervised autonomous mode — responses deployed automatically or manually per configuration and countermeasure type — and, being microservice-based, the platform adds sensors, effectors, and functions without altering the core, spanning fixed and vehicle-mounted configurations alike. Figure 3 shows the platform fusing radio-frequency detection with electro-optical tracking of the same target while displaying sensor and effector coverage on the map.

Each SUDA phase carries measured figures, summarised in Table 1 — engineering values read from the system's own instrumentation and logs during commissioning and routine operation, quoted to characterize the architecture, not as contractual guarantees. The Sense layer illustrates the data-volume problem concretely: a thermal panoramic sensor of 1280 × 1024 pixels, driven by a high-end video-

processing platform developed at the Institute [18], covers the full circle in 36 sectors every 2.3 s. The Understand layer illustrates the compression AI achieves: the raw flow of about fifty candidate detections per second, dominated by clouds and birds, is reduced to approximately one plausible target per second by tracking, cloud and bird filtering, and direct referral of close targets to AI classification — techniques building on the Institute's robust tracking and learned detection in thermal and short-wave-infrared imagery [19], [20], [21] — a fifty-to-one reduction before any human sees the picture: the difference between an operator who is overwhelmed and one who is informed.

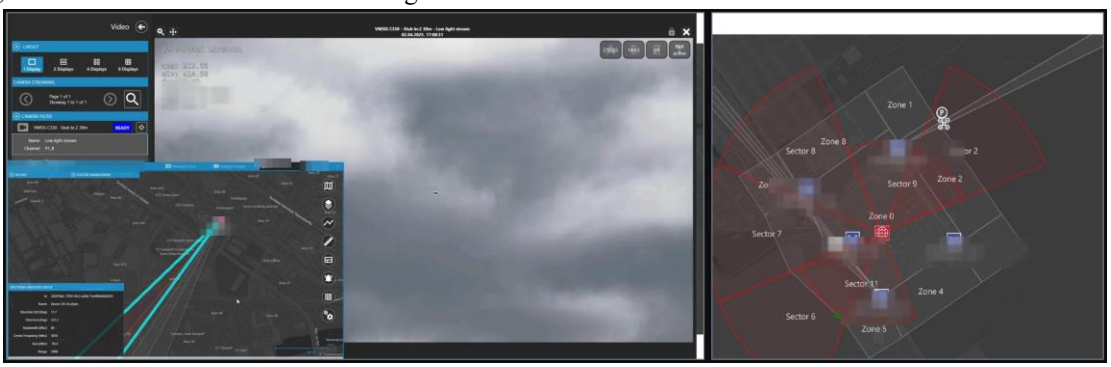


Figure 3. The ORAO platform in operation at a demonstration site: the live optical stream with telemetry, the map with bearing lines from RF spectrum sensing, and the zone and sector overlay of sensors and effectors on a single console (identifying details redacted)

Table 1. SUDA phases instantiated on the ORAO platform, with characteristic measured quantities

SUDA phase	Function	Realisation	Characteristic figures
Sense	Multi-sensor acquisition, conditioning, time reference	Thermal panoramic head; RF spectrum sensing; vMSIS3 EO/IR turret with laser rangefinder; a dedicated passive acoustic sensor; X-band radar where permitted	Full circle in 2.3 s; RF spectrum 20 MHz–6 GHz
Understand	Track fusion, COP, AI classification	ARES Edge and ARES Core on current-generation GPU; classes: drone, bird, aircraft, helicopter; ontology layer; Kalman filtering with deep-learning re-identification	Above 100 TOPS; inference latency below 100 ms; four parallel video streams; 50 raw detections/s reduced to about 1/s
Decide	Threat evaluation, option preparation, ROE constraints	ORAO decision services with ARES Strategic: AI and rule-based events and actions; LLM-generated COA; supervised autonomous operation	Operator authority by default; automatic response only where explicitly configured
Act	Effector and sensor tasking, monitoring, feedback	Cueing of the vMSIS3 turret; RF jamming with selective radiation and coverage control; effect assessment returned to the COP	Turret classifies a cued object in under 2 s from any previous position

The Decide and Act layers are equally concrete. Rules of engagement appear as configuration: the protected installation applies selective control of jammer radiation power and coverage zones so that countermeasures do not reach the surrounding populated area — an ROE constraint compiled into system behavior exactly as Section 3.3 requires — and safeguards prevent countermeasures from triggering without authorization unless a fully automatic mode has been explicitly enabled. The Act layer closes the loop within the electro-optical turret's two-second verification window, automatically directed to coordinates supplied by another sensor and returning a classified, ranged, and geo-located object to the picture that cued it; the Institute has reported this turret's optical platform and pointing performance [17].

5. HUMAN–MACHINE TEAMING

Whether AI-assisted C2 delivers value depends less on model accuracy than on the quality of the human–machine team; three requirements structure the problem. The first is a deliberate allocation of function. The taxonomy of Parasuraman, Sheridan and Wickens [24] distinguishes automation of information acquisition, analysis, decision selection and action implementation, each at graded levels. In the SUDA framing, high automation is appropriate in Sense and much of Understand, where machines outperform humans in throughput and vigilance — the fifty-to-one reduction of Section 4 is no human task at all. In Decide, automation prepares and ranks options, but authority over the use of force remains with the human commander; delegation beyond that point must be explicit, bounded, and reversible—the standard allied policy now codifies autonomy in weapon systems [25]. Autonomy is not fixed: control tightens when time permits, and narrowly defined responses are pre-authorized where reaction times make deliberation impossible — supervised autonomous operation in ORAO is precisely a graduated setting of this parameter.

The second is explainability. A recommendation that cannot be explained will in practice be either ignored or blindly followed — both failure modes. The field systematized in the DARPA XAI program [26] offers

many techniques, but the operational requirement is specific: explanations in the operator's terms — which evidence, from which sensors, drove this ranking; what the confidence is; what would change the conclusion — delivered at the moment of decision and honest about uncertainty, the transparency recent work identifies as the foundation of operator situation awareness in human-AI teams [27]. A classification shown on the console with its confidence value alongside the imagery behind it is a modest but real instance of the principle.

The third is trust calibration: over-trust produces automation bias — operators accepting machine output uncritically, a failure mode documented across military decision-support deployments [28] — while under-trust wastes the machine's contribution. Calibration is built through transparency of confidence, training in the system's characteristic strengths and failure modes, and consistent logging. Responsibility is not transferred to software: the commander remains accountable, and the system must be designed so that this accountability is real, traceable, and fair.

6. IMPLEMENTATION CONSIDERATIONS

6.1. Edge and central inference

A practical architecture distributes intelligence between the sensor edge and the center: latency-critical inference (detection, first-line classification, data conditioning) at the edge, where it also relieves communication links; resource-intensive functions (model training, large-scale fusion, long-horizon analytics) centrally. In the deployed system, all AI processing occurs locally, in the field or vehicle, so the system remains operational without a network connection, and the edge platform is retrained locally on newly collected data. The edge–center connection assumes contested communications: edge nodes continue autonomously on locally cached models when links fail, the COP degrades gracefully, and synchronization resumes when connectivity returns. Adoption in an existing C2 system is likewise incremental: the layer first enters as an advisory overlay through standardized interfaces — including NATO STANAG and legacy national formats — and services migrate to integrated status as confidence and accreditation mature, building the doctrinal and trust foundations of Section 5.

6.2. Security and sovereignty

AI-assisted C2 operates on data whose confidentiality, integrity and availability are matters of national security. Data in transit and at rest are therefore protected by cryptographic means under national control — centralized key management, an interoperable public-key infrastructure and secure key distribution [29]; the chain's trustworthiness rests on key generation, hence hardware true-random generation as the foundation — the Institute's research on autonomous symmetric ciphering and key distillation [22], [23] addresses exactly this chain. The C2 network also defends itself with the same family of methods it offers the commander: machine-learning detection of anomalies in the network's own behavior [30]. Models are assets: provenance known, training data controlled, updating performed within the national security perimeter — models trained on classified operational data cannot depend on external infrastructure — the deeper reason why local retraining at the edge is an architectural requirement, not merely a latency optimization. Emission security is addressed through hardware design and accreditation. Sovereign control over data, models and cryptography is thus a precondition of operational trust.

7. CYCLE-TIME ANALYSIS OF AN ENGAGEMENT

Consider a protected critical-infrastructure site of the kind described in Section 4 — roughly one square kilometer, observed by a thermal panoramic head, RF spectrum sensing and an agile electro-optical turret, with a countermeasure tower under selective radiation control, integrated through ORAO; the objective is to stop small drones before they approach within 500 m of the perimeter. In Sense, the panoramic head sweeps the full circle every 2.3 s while the RF layer listens across the spectrum; a drone entering the outer zone is registered by one or both — an emitter by its control link, a silent one by its thermal signature against the sky. In Understand, the raw flow is compressed before an operator sees anything: detections are bound into tracks, cloud and bird returns rejected, and the surviving track fused with the RF bearing into one object in the COP, carrying position, classification, confidence, and contributing sensors. ORAO cues the turret; its edge-AI returns a class — drone rather than bird, aircraft or helicopter — with a confidence value in under two seconds, while the laser rangefinder converts bearing and range into coordinates: a few seconds after first detection, the operator has one identified, located, ranked object instead of fifty ambiguous stimuli per second.

In Decide, the platform prepares the response: which countermeasure applies, which zone it may radiate into, whether the rules permit automatic action or require approval—and, because the ROE are compiled into the configuration, it offers no option that would radiate into the prohibited sector. The operator, alerted by a single track rather than a continuous stream, reviews the classification and supporting imagery, then authorizes the measure. In Act, the jammer radiates its selective pattern, the turret observes, and the outcome — the drone descending, returning or landing — is recorded as effect assessment, refining both the picture and the models. The same discipline scales beyond soft-kill: with layered gun-based and directed-energy effectors from the Institute's portfolio, the chain from first radar contact with a multi-drone raid at 8 km to confirmed neutralization closes within about 45 seconds.

The interval from first stimulus to a directed, compliant response is thus compressed from the timescale of manual correlation across separate consoles to that of a supported decision, with every step transparent, logged and under human authority; and it decomposes — 2.3 s of panoramic revisit in Sense; sub-100 ms inference and under 2 s of turret verification in Understand; a human decision bounded by how well the option was prepared; a tasking latency in Act — each term a separate engineering problem with a separate owner. That is what it means to say that

SUDA turns cycle time from a metaphor into a measurable system property.

8. CONCLUSION

This paper has presented an architectural framework for AI-assisted decision support in C2 systems, organized around the SUDA cycle and demonstrated on a system in operational service. Its central claim is that decision advantage arises not from replacing the human decision-maker but from restructuring the C2 information flow so that machines do what they do best — acquire, fuse, classify, detect anomalies, prepare and rank options — while human judgement is applied, deliberately and accountably, where it matters most. The Vlatacom C7ISTAR realisation, in which the ORAO platform carries the ARES three-tier AI core, substantiates the claim phase by phase with the measured figures of Table 1. Together those figures describe a single object — a short, instrumented and accountable path from sensor to decision-maker and from decision to effect — and it is that path, not any individual algorithm, that modern command support has to engineer. The sovereign, end-to-end integration of sensors, intelligence and effectors within one platform is becoming the defining requirement of the field.

Future work follows three lines: standard metrics and instrumentation for per-phase cycle-time measurement across heterogeneous systems; test, evaluation and accreditation of learning-based functions, including behaviour under distribution shift and adversarial conditions; and interoperability of SUDA-structured decision support across systems and coalition boundaries.

Acknowledgement

The research underlying the sensing and AI components of this paper was funded in part by the Vlatacom Institute of High Technologies (project no. 157, VMSIS3_Advanced) and by the Science Fund of the Republic of Serbia (Grant no. 6524745, AI-DECIDE).

References

[1] BOYD,J.R.: A Discourse on Winning and Losing (ed. Hammond,G.T.), Air University Press, Maxwell AFB, 2018.

[2] BREHMER,B.: The Dynamic OODA Loop: Amalgamating Boyd's OODA Loop and the Cybernetic Approach to Command and Control, Proc. 10th ICCRTS, McLean, VA, 2005.

[3] ENDSLEY,M.R.: Toward a Theory of Situation Awareness in Dynamic Systems, Human Factors, 37(1) (1995) 32–64.

[4] VULIĆ,I.: Mobilni geografski informacioni sistemi u sistemima C4I2 (Mobile Geographical Information Systems in C4I2 Systems), Vojnotehnički glasnik, 54(3) (2006) 354–367.

[5] VULIĆ,I.: Komandno-informacioni sistemi (Command-Information Systems), University of Defence in Belgrade, Belgrade, 2026 (in Serbian, in press).

[6] JOHNSON,J.: Automating the OODA Loop in the Age of Intelligent Machines: Reaffirming the Role of Humans in Command-and-Control Decision-Making in the Digital Age, Defence Studies, 23(1) (2023) 43–67.

[7] STEINBERG,A.N., BOWMAN,C.L., WHITE,F.E.: Revisions to the JDL Data Fusion Model, Proc. SPIE 3719, 1999, pp.430–441.

[8] U.S. DEPARTMENT OF DEFENSE: Summary of the Joint All-Domain Command and Control (JADC2) Strategy, Washington, D.C., 2022.

[9] VULIĆ,I., PRODANOVIĆ,R., TOT,I., BOGIĆEVIĆ,D.: Model for Authenticating the Internet of Military Things and Internet of Battlefield, Proc. ICIST 2020, Kopaonik, 2020, pp.206–210.

[10] HALL,D.L., LLINAS,J.: An Introduction to Multisensor Data Fusion, Proceedings of the IEEE, 85(1) (1997) 6–23.

[11] MOSTAFA,M.A.R.M., VUČETIĆ,M., STOJKOVIĆ,N., LEKIĆ,N., MAKAROV,A.: Fuzzy Functional Dependencies as a Method of Choice for Fusion of AIS and OTHR Data, in: Smart Sensors and Devices in Artificial Intelligence, Springer, Basel, 2022.

[12] VULIĆ,I., BORISOV,M., PRODANOVIĆ,R., RANČIĆ,D., PETROVIĆ,V.M., STANKOVSKI,S., OSTOJIĆ,G.: Protection of Digital Elevation Model — One Approach, Applied Sciences, 12(19) (2022) 9898.

[13] GOECKS,V.G., WAYTOWICH,N.: COA-GPT: Generative Pre-trained Transformers for Accelerated Course of Action Development in Military Operations, arXiv:2402.01786, 2024.

[14] RIVERA,J.-P., MUKOBI,G., REUEL,A., LAMPARTH,M., SMITH,C., SCHNEIDER,J.: Escalation Risks from Language Models in Military and Diplomatic Decision-Making, Proc. ACM FAccT '24, Rio de Janeiro, 2024.

[15] PARK,S., KIM,H.T., LEE,S., JOO,H., KIM,H.: Survey on Anti-Drone Systems: Components, Designs, and Challenges, IEEE Access, 9 (2021) 42635–42659.

[16] SEMENYUK,V., KURMASHEV,I., LUPIDI,A., et al.: Advances in UAV Detection: Integrating Multi-Sensor Systems and AI for Enhanced Accuracy and Efficiency, Int. J. Critical Infrastructure Protection, 49 (2025) 100744.

[17] LATINOVIĆ,N., POPADIĆ,I., TOMIĆ,B., SIMIĆ,A., MILANOVIĆ,P., NIJEMČEVIĆ,S., PERIĆ,M., VEINOVIĆ,M.: Signal Processing Platform for Long-Range Multi-Spectral Electro-Optical Systems, Sensors, 22(3) (2022) 1294.

[18] POPADIĆ,I., PERIĆ,M., TOMIĆ,B., SIMIĆ,A., MILANOVIĆ,P.: High-End Video Processing Platform for Electro-Optical Systems, Proc. OTEH 2020, Belgrade, 2020.

[19] PAVLOVIĆ,M., BANJAC,Z., KOVAČEVIĆ,B.: Object Tracking in SWIR Imaging Based on Both Correlation and Robust Kalman Filters, IEEE Access, 11 (2023) 63834–63851.

[20] LATINOVIĆ,N., VUKOVIĆ,T., PETROVIĆ,R., PAVLOVIĆ,M., KADIJEVIĆ,M., POPADIĆ,I., VEINOVIĆ,M.: Implementation Challenge and

Analysis of Thermal Image Degradation on R-CNN Face Detection, Telfor Journal, 12(2) (2020) 98–103.

[21] PAVLOVIĆ,M., MILANOVIĆ,P., STANKOVIĆ,M., PERIĆ,D., POPADIĆ,I., PERIĆ,M.: Deep Learning Based SWIR Object Detection in Long-Range Surveillance Systems: An Automated Cross-Spectral Approach, Sensors, 22(7) (2022) 2562.

[22] MILOSAVLJEVIĆ,M., RADOMIROVIĆ,J., UNKAŠEVIĆ,T., BOŽILOVIĆ,B.: One Class of Ideally Secret Autonomous Symmetric Ciphering Systems Based on Wiretap Polar Codes, Mathematics, 12(23) (2024) 3724.

[23] RADOMIROVIĆ,J., MILOSAVLJEVIĆ,M., BANJAC,Z., JOVANOVIĆ,M.: Secret Key Distillation with Speech Input and Deep Neural Network-Controlled Privacy Amplification, Mathematics, 11(6) (2023) 1524.

[24] PARASURAMAN,R., SHERIDAN,T.B., WICKENS,C.D.: A Model for Types and Levels of Human Interaction with Automation, IEEE Trans. SMC — Part A, 30(3) (2000) 286–297.

[25] U.S. DEPARTMENT OF DEFENSE: DoD Directive 3000.09, Autonomy in Weapon Systems, Washington, D.C., 2023.

[26] GUNNING,D., AHA,D.W.: DARPA's Explainable Artificial Intelligence (XAI) Program, AI Magazine, 40(2) (2019) 44–58.

[27] ENDSLEY,M.R.: Supporting Human-AI Teams: Transparency, Explainability, and Situation Awareness, Computers in Human Behavior, 140 (2023) 107574.

[28] PROBASCO,E., TONER,H., BURTELL,M., RUDNER,T.G.J.: AI for Military Decision-Making: Harnessing the Advantages and Avoiding the Risks, CSET, Georgetown University, 2025.

[29] PRODANOVIĆ,R., VULIĆ,I.: Model za PKI interoperabilnost u Republici Srbiji (A Model for PKI Interoperability in Serbia), Vojnotehnički glasnik, 65(2) (2017) 530–549.

[30] PROTIĆ,D., STANKOVIĆ,M., PRODANOVIĆ,R., VULIĆ,I., et al.: Numerical Feature Selection and Hyperbolic Tangent Feature Scaling in Machine Learning-Based Detection of Anomalies in the Computer Network Behavior, Electronics, 12(19) (2023) 4158.