COST-EFFECTIVE MILITARY ISR VIA SMALL MULTI-SENSOR UAS: LEVERAGING MINIATURISATION FOR ENHANCED CAPABILITY AND AGILE DEVELOPMENT

Wilken. Hemme*, Jens. Halbig*, Benjamin. Kormann*
* Hochschule München, Faculty 04, Lothstr. 64, Munich, Germany

Abstract

The convergence of miniaturisation and advanced on-board compute is enabling a new class of cost-effective, multi-sensor Unmanned Aerial System (UAS) that are reshaping the modern battlespace. However, unlocking their full potential requires a shift from remotely piloted data collectors to truly autonomous, intelligent and coordinated teams. This paper introduces a novel, three-layer cognitive architecture that integrates a utility-based system for strategic prioritisation, a goal oriented action planner for tactical flexibility, and a stigmergic layer for decentralised coordination. Validation in a simulated military Intelligence, Surveillance, and Reconnaissance (ISR), attack and supply mission demonstrates that the stigmergic coordination layer yields a profound increase in efficiency, reducing mission completion time by 46.5% and total distance travelled by 47.7%. The framework's data-driven design also facilitates the agile development of tactics, allowing new mission sets and behaviours to be rapidly configured. The results validate a practical path towards scalable, intelligent autonomous systems capable of emergent, coordinated action in dynamic environments.

Keywords

UAS; Military ISR; Autonomous Systems; Hybrid AI Architecture; Emergent Behaviour; Goal-Oriented Action Planning (GOAP); Utility System; Stigmergy

1. INTRODUCTION: THE NEW STRATEGIC PILLAR OF UNMANNED SYSTEMS

1.1. The Miniaturisation Revolution and the Rise of Asymmetric Warfare

We are currently witnessing first-hand how the strategic landscape of modern conflicts is being fundamentally reshaped. Powerful sensor and processing capabilities that can be integrated into small, low-cost unmanned platforms [1] are already significantly changing the battlefield. This technological shift has led to a profound asymmetric advantage, as almost anyone now has access to advanced ISR and attack capabilities. These low-cost systems are now in a position to effectively threaten and neutralise high-value military assets [2, 3]. The proven success of these platforms in recent conflicts underscores the urgent need to understand and master the autonomous warfare domain [4].

1.2. From On-Board Compute to Enhanced Cognitive Capability

This revolution is primarily driven by the increased performance of smaller hardware, which makes improved functions possible in the first place. The current focus on Artificial Intelligence (AI) has triggered an unprecedented surge of innovation in the research and development of these chips, with the result that computing power related to AI doubles approximately every six months [5]. This exponential growth is transforming UAS from simple flying cameras into powerful *flying brains*, embodied by platforms such as the one developed for our work (Figure 1).

FIG 1. The physical multi-sensor UAS platform used in our flight tests. The integration of advanced sensors into such a compact, low-cost airframe highlights the central challenge addressed by this work: the need for a sophisticated cognitive architecture to transform raw data into intelligent, autonomous actions.

A multitude of groundbreaking developments has also made the basic operational capability of UAS possible in the first place. Decades of fundamental developments in the areas of integrated circuits [6], energy storage [7] and robust communication [8] are just a few significant advances. Our collected multimodal sensor data (Figure 2), gathered during flight tests with this platform [9, 10], reveals a new challenge: it is not just a matter of collecting

CC BY 4.0

doi: 10.25967/650406

data, but also of **interpreting** it and making **intelligent decisions** based on it.
This is a crucial skill that UAS must master. In the face of advanced electronic warfare that can disrupt communication links, the platform is forced to complete its mission entirely on its own. The need for robust on-board autonomy in such communication-contested environments means the system cannot rely on a human operator for continuous guidance [11].

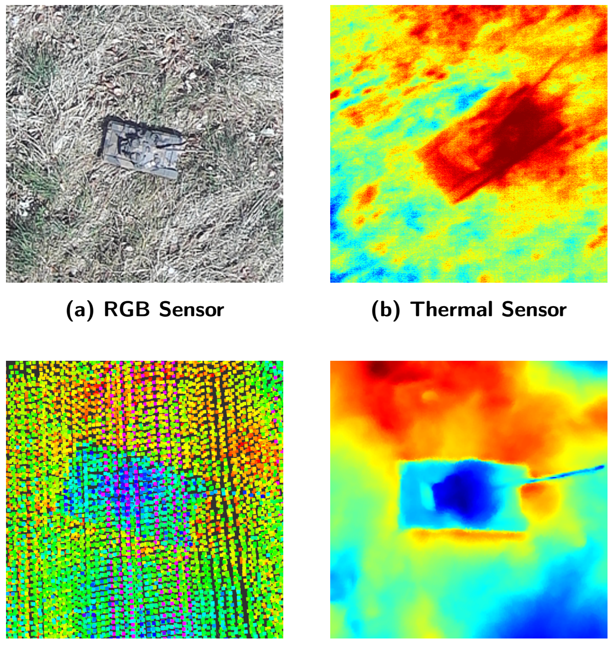

(a) RGB Sensor **(b) Thermal Sensor**

(c) LiDAR Sensor **(d) Depth Sensor**

FIG 2. Multi-modal sensor data of a military tank model, captured during real-world flight tests. This demonstrates the rich data available from a single miniaturised platform, highlighting the cognitive challenge of autonomous interpretation needed for *enhanced capability*.

1.3. The Need for Agile Development of Autonomous Tactics

To exploit this potential, agility is required in two different but interrelated forms. On the one hand, agile competence development is necessary at the strategic level. Operators and engineers must be able to quickly define, test and deploy new mission sets and autonomous behaviours. In the new world, it is essential that this is not held back by lengthy, monolithic software development cycles, because the battlefield changes within hours, not months, and our systems must keep pace.
Secondly, and more importantly, the autonomous system itself must possess *agile operational capabilities*. Enables dynamic redistribution of tasks during flight without intervention from the ground, reducing the cognitive load on the remote operator. Once deployed, the team of Unmanned Aerial Vehicles (UAVs) must be able to dynamically create and adapt its own tactical plans. In a constantly changing, unpredictable and hostile environment, this ability is essential, as the *flying brain* cannot rely on a pre-planned script. It must observe, orient itself, decide and act, effectively developing its own solutions to problems as they arise. This article presents an architecture that offers both forms of agility and bridges the gap between high-level mission objectives and effective, coordinated action.

2. A HYBRID ARCHITECTURE FOR EMERGENT AUTONOMY

The solution proposed is a modular, three-layer cognitive architecture (Figure 3) designed to solve three distinct challenges of autonomous multi-agent missions: strategic prioritisation, tactical planning, and decentralised coordination. Its design, with roots in the computer game industry, is founded on the concept of goal-oriented agents whose collective behaviour emerges from local interactions rather than a central plan.

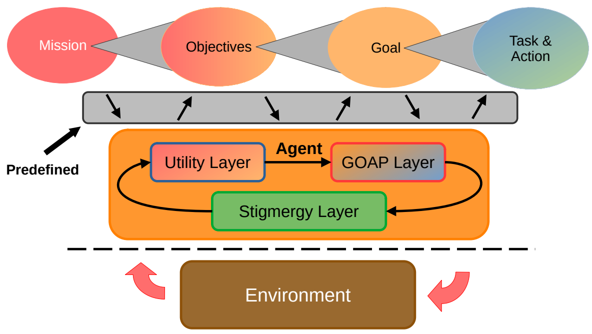

FIG 3. The proposed three-layer hybrid architecture, illustrating the interaction between the strategic, tactical, and coordination layers.

2.1. Architectural Layers

The architecture integrates three specialised components that operate in a continuous cycle:
Utility System (UTS) forms the strategic layer, tasked with answering *what* an agent should do. With its roots in decision theory, it continually evaluates potential objectives based on a set of weighted criteria known as **considerations**. This allows the system to fuse heterogeneous factors—such as target proximity and threat level—into a single utility score, ensuring the agent pursues the most rational goal or objective at any given moment [12, 13].
Goal-Oriented Action Planning (GOAP) serves as the tactical layer to determine *how* a selected goal can be achieved. In contrast to a top-down planner like an Hierarchical Task Network (HTN), GOAP works bottom-up, finding novel sequences of atomic actions to achieve a goal. This provides significant flexibility, enabling an agent to generate its own solutions to unforeseen problems rather than following a rigid predefined script [14].
Stigmergy provides the coordination layer, answering *which* agent should pursue a given task. By leaving a digital mark in the shared `WorldState`, an actor signals their intention. Other team members can take this mark into account and incorporate it into their own utility calculations. This simple mechanism of indirect com-

CC BY 4.0

munication leads to emergent deconfliction and an efficient, decentralised division of labour, which is ideal for communication-contested environments [15].

2.2. Agile Design via Data-Driven Configuration

Crucially, the entire system is realised through a data-driven approach. The agent's core cognitive structure—from its `WorldState` representation to its domain-specific logic—is defined in human-readable YAML Ain't Markup Language (YAML) files. This provides significant architectural agility, enabling mission designers to rapidly reconfigure and adapt the system's capabilities for new operational contexts without modifying the core software. This agile design at the system level provides the foundation for the operational agility required during a mission.

3. CASE STUDY: DEMONSTRATING EMERGENT GOAL FORMATION IN A SIMULATED ATTACK AND MEDICAL RESUPPLY MISSION

To validate the architecture, a special simulated military scenario was carried out based on the targets and environment of our real flight tests. The scenarios served not only to test the implementation, but more importantly to evaluate the concrete benefits, such as efficiency, adaptability and coordination, of the core principles of the architecture. Crucially, these tests are intended to demonstrate the agent's ability to develop agile tactics and effective collective teams. The aim is to prove that the enhanced capabilities promised in the title are not merely theoretical, but also materialise in practice.

3.1. Simulation Setup

To validate the architecture and its ability to generate emergent collective behaviour, a complex military scenario was configured. This configuration forms the experimental environment for the evaluation described in the following sections. The entire mission is defined via external YAML configuration files to emphasise the domain-independent nature and principles of *agile design* of the core architecture.
As shown in Figure 4, the simulation involves a hostile operational area in which a team of three UAVs must carry out a combined attack and supply soldiers with medical equipment. The initial `WorldState` creates a map with numerous dynamic threats, including anti-aircraft vehicles, tanks, soldiers, and signal reconnaissance vehicles, which pose a threat to the agent team. The main objective is to supply the soldiers deep in enemy territory with medical equipment. To achieve this, the UAVs must first eliminate threats such as tanks, anti-aircraft defences and signal reconnaissance vehicles in order to safely deliver medical supplies to soldiers.

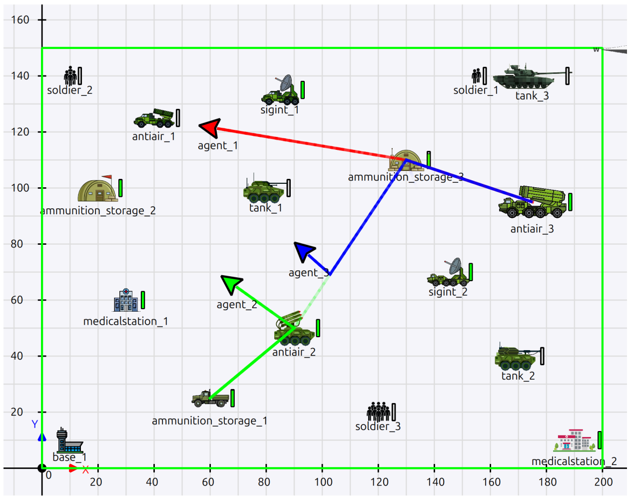

FIG 4. **Simulation environment of a military scenario. Three UAVs carry out an attack and a Medical Resupply Mission. The objective is to evacuate soldiers in an area surrounded by anti-aircraft defences, tanks and radar reconnaissance vehicles**

3.2. Scenario 1: Cost-Effectiveness and Emergent Coordination

To measure how much the stigmergic layer helps to improve the efficiency and coordination of the agents, the mission was run 25 times, both with and without the mechanism turned on. The results, shown in Table 1, show a clear improvement in collective efficiency when agents can coordinate indirectly.

TAB 1. **Performance comparison for military Scenario with and without the stigmergy mechanism enabled (averaged over 25 runs)**

Metric	Stigmergy (off)	Stigmergy (on)
Mission Completion Time (s)	145.85 ± 4.84	$\mathbf{78.12 \pm 3.69}$
Total Distance Traveled (Units[1])	10047.14 ± 0.00	$\mathbf{5254.12 \pm 138.30}$
Avg. GOAP Planning Time (ms)	24.21 ± 28.78	14.04 ± 20.86
Avg. UTS Decision Time (μs)	1.39 ± 1.24	1.34 ± 0.90

When the stigmergy layer is deactivated, the agents act completely egocentrically. Lacking shared situational awareness, they often pursue the same high-utility objectives, leading to redundant travel and inefficient task allocation. This baseline performance resulted in an average mission completion time of **145.9** seconds and a total distance travelled of over **10,000** units.
Activating the stigmergy level fundamentally changes this dynamic. By leaving digital markers indicating their current goals, the agents create a shared, coordinated intention to solve the mission. This enables spontaneous

[1] Arbitrary spatial units within the simulation grid.

CC BY 4.0

conflict avoidance during tasks and a more logical distribution of resources within the team. The effects are dramatic: the time to complete the mission was reduced by **46.5%** to just **78.1** seconds, and the total distance travelled was reduced by **47.7%**. This highlights a significant increase in operational speed and cost efficiency, as the same goals are achieved with significantly less time and energy.

Interestingly, the average GOAP planning time also decreased from **24.2** ms to **14.0** ms. This is attributed to the stigmergic layer guiding agents towards distinct, deconflicted goals. Replanning from these more tactically advantageous positions fundamentally changes the A* search, as the cost of actions leading to the new goal can be significantly lower than for all alternatives. This allows the planner to prune unfeasible branches more aggressively, reducing the search space and leading to faster convergence.

Meanwhile, the UTS decision time remained negligible at approximately **1.3** μs in both configurations, confirming that the stigmergic layer adds no significant cognitive overhead to the strategic decision-making process.

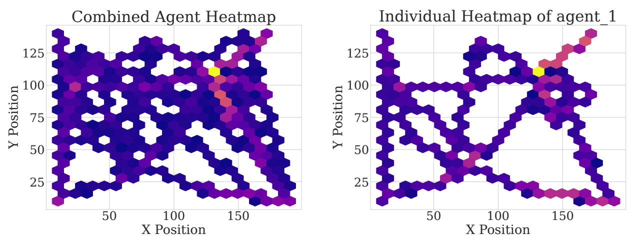

(a) Combined Heatmap **(b) Agent 1 Heatmap**

(c) Agent 2 Heatmap **(d) Agent 3 Heatmap**

FIG 5. Influence of the stigmergy layer on the spatial distribution of agents. The individual agent heatmaps (b-d) reveal an emergent spatial deconfliction, with each agent focusing on a distinct sector of the operational area.

The stigmergic layer's influence is also evident in the emergent spatial distribution of the agents, which is visually confirmed by the heatmaps in Figure 5. The maps reveal a clear division of labour: Agent 1 focuses on the lower-right, Agent 2 on the upper-left, and Agent 3 covers the central corridor. This emergent deconfliction ensures more comprehensive area coverage with less redundant movement, directly contributing to the efficiency gains reported in Table 1. Without stigmergy, this partitioning does not occur, as all agents independently pursue the same initial high-utility targets.

3.3. Scenario 2: Agile Re-tasking against Dynamic Threats

This second scenario serves to validate one of the core statements of the architecture with regard to operational agility. Specifically, it tests the UTS's ability to dynamically re-prioritise goals in response to an emergent threat by comparing a baseline mission execution against a variation where the tactical situation is dynamically altered. In the baseline configuration, with no immediate threats, the team of UAV's adheres to the pre-defined mission priorities. They first neutralise the signal intelligence vehicles (*sigint_1*, *sigint_2*), followed by the anti-air assets, the remaining tanks, and finally deliver medical supplies to the soldiers before returning to base. The resulting goal execution timeline, shown in Figure 6, provides a predictable performance benchmark.

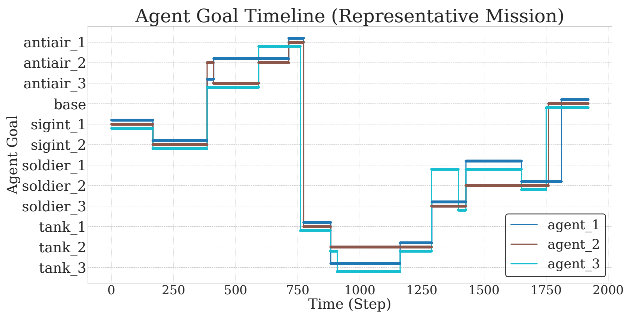

FIG 6. Goal execution timeline for the baseline scenario without dynamic threats. The team of UAV systematically selects the goals according to their static priority.

The second configuration introduces a critical dynamic element: *tank_3* is positioned in close proximity to *soldier_1*. While the static priority of tanks remains low, the UTS contains a consideration that heavily penalises threats near high-value friendly assets.

As shown in Figure 7, this contextual factor causes the team to dynamically override the default priority. Immediately after engaging the *SIGINT* targets, the agents redirect to neutralise *tank_3*—an action taken far earlier than its static priority would dictate. Once this immediate threat to the soldiers is eliminated, the team resumes the standard mission profile, clearing the remaining threats before proceeding with the medical resupply.

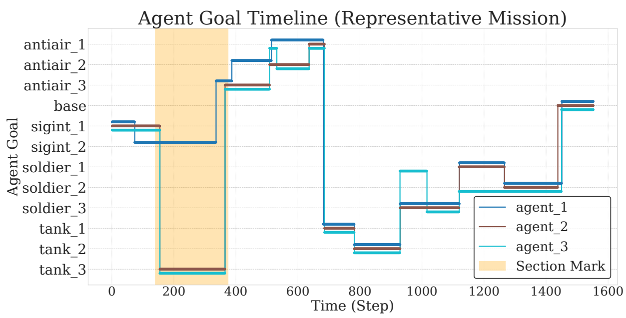

FIG 7. Goal execution timeline with a dynamic threat. The UTS overrides the static priority to engage *tank_3* early, demonstrating agile re-tasking.

CC BY 4.0

This result is a crucial validation of the architecture's claim to *agile operational capability*. It demonstrates that the UTS is not merely executing a static script but is performing true, context-aware reasoning. By weighing multiple factors—static priority, proximity, and threat level—it makes an intelligent, emergent decision that a rigidly planned system could not. This ensures the team responds effectively to the most critical aspects of the evolving battlespace.

3.4. Discussion: Architectural Performance and Tactical Implications

An analysis of the architectural components reveals the performance trade-offs inherent to this hybrid design, which are the foundation of its *enhanced capability*. This capability is driven by the direct interplay between the reactive UTS and the deliberative GOAP planner.

The UTS forms the strategic core, enabling the agile retasking demonstrated in Scenario 2. With an average decision time of approximately **1.3** µs, it provides the near-instantaneous prioritisation required for real-time responsiveness in a dynamic military scenario. This efficiency ensures the team can adapt its high-level objectives and goals without incurring significant cognitive overhead.

In contrast, the GOAP planner provides tactical depth, generating the sequence of actions required to achieve a selected goal. However, this deliberative capability comes at a computational cost. Since the implementation utilises an A* search algorithm at its core, the time required to find a solution scales with the complexity of the problem. This is evident in Figure 8, which plots planning time against plan length across all simulation runs. The data reveals a clear exponential growth trend (with $R^2 = 0.820$), highlighting a key computational limitation of the architecture.

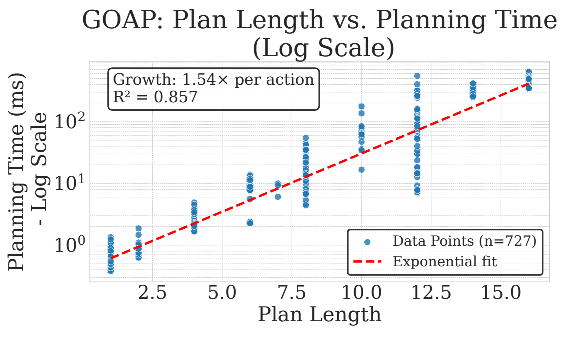

FIG 8. The relationship between GOAP plan length and planning time, averaged over all scenarios. The data shows a clear exponential growth trend, highlighting a key computational limitation for missions requiring long action sequences.

The stigmergic layer acts as the connective tissue between these two components at a multi-agent level, enabling the emergent coordination seen in Scenario 1. It provides the essential collective context that feeds into each agent's individual decision cycle. The digital markers left by team members directly inform the UTS, allowing it to de-prioritise objectives that are already being addressed. This prevents the costly redundant behaviour of uncoordinated agents and is the primary driver of the **46.5%** reduction in mission time.

This three-part architecture creates a powerful performance dynamic. Stigmergy provides the collective context with minimal overhead; the UTS leverages this context to select the most critical objective with near-instantaneous speed, whilst the GOAP provides the tactical intelligence to generate a plan to achieve it. Although the planner proved to be sufficiently powerful for the tested scenarios, its exponential nature represents a bottleneck in missions that require longer or more complex sequences of actions. This trade-off is central to the design of the system. Future work could mitigate this limitation by investigating optimisations to the planner, such as implementing a regressive search or migrating the planner to a more powerful language such as *C++*.

3.4.1. Performance Comparison with a State-of-the-Art Baseline

To contextualise these findings, a comparative analysis was performed against *Fast Downward*, a state-of-the-art classical planner [16]. Both planners solved the same set of problems using optimal A* search configurations. Given the presence of outliers, the **median** is used as a more robust measure of typical performance.

The results, summarised in Table 2 and Figure 9, confirm that *Fast Downward* is substantially more performant. This is an expected outcome, given its highly optimised *C++* implementation compared to our *Python* prototype.

TAB 2. Key performance metrics of the implemented GOAP (914 runs) (Max planning length 12) vs. Fast Downward (Max planning length 36) (200 runs), using the median to represent typical performance.

Metric	Our GOAP	Fast Downward
Median Planning Time (ms)	8.13	**1.38**
R^2 (Exponential Fit)	0.875	**0.838**

However, this performance benchmark does not capture a crucial feature unsupported by classical planners: the ability to perform numerical comparisons within the planning process. Our GOAP implementation can define preconditions that dynamically compare `WorldState` values, such as *agent.ammunition* $>$ *target.health*.

This capability is essential for the agile behaviour demonstrated in Scenario 2, as a classical planner cannot natively reason about such numerical relationships. This establishes a critical trade-off: whilst *Fast Downward* excels in raw speed for purely symbolic domains, our GOAP implementation provides the flexibility required for environments that demand dynamic, numerical reasoning at the action level.

CC BY 4.0

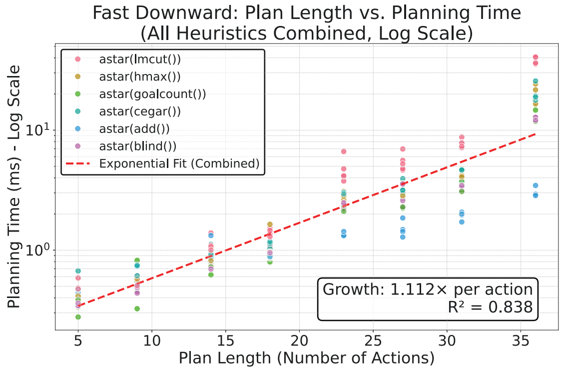

FIG 9. Planning time growth for the *Fast Downward* baseline. The data confirms the expected exponential trend for an optimal A* search, while demonstrating substantially faster performance compared to the GOAP planner.

4. CONCLUSION: A FRAMEWORK FOR AGILE CAPABILITY DEVELOPMENT

4.1. Summary of Contributions

This work has delivered on the promise of its title by presenting a concrete architectural solution for a new generation of autonomous systems. We have demonstrated a framework that provides *enhanced cognitive capabilities* and enables the *agile development* of both system-level mission sets and on-the-fly operational tactics.

Compared to established alternatives, the proposed hybrid architecture offers a robust, goal-oriented adaptability often lacking in purely reactive systems like Finite State Machines or Behaviour Trees. Furthermore, its stigmergic layer enables effective team coordination without the high communication overhead and single point of failure associated with centralised strategies, such as market-based systems. The results confirm that by integrating a UTS, GOAP, and a stigmergic layer, concepts from the gaming industry can be effectively adapted to create sophisticated, emergent behaviours in complex autonomous systems.

4.2. Threats to Validity

Whilst the results are promising, two limitations must be acknowledged. Firstly, **internal validity** is constrained by the simulation environment. The agent planners were executed sequentially to avoid race conditions, and real-world complexities like sensor noise and communication latency were abstracted away. The observed performance may therefore be optimistic. Secondly, **external validity** is limited as the architecture was tested only in this military domain. A direct performance comparison against other established planners was beyond the scope of this work, and conclusive evidence of robustness requires deployment onto a physical system.

4.3. Future Work: Towards Rapid and Agile Capability Deployment

The research presented here serves as a robust foundation for future work aimed at closing the simulation-to-reality gap. The primary objective is the deployment of this architecture onto physical multi-UAS platforms to validate its performance in real-world conditions. Future efforts will also focus on streamlining the pipeline for *agile capability development*. The ultimate goal is to create high-level tools that would allow an operator, with no programming expertise, to rapidly define, simulate, and deploy entirely new autonomous behaviours for the team. This would represent the final step: making this powerful technology truly accessible and adaptable at the speed of relevance, allowing tactical innovation to occur in the field, not just in the lab.

5. LIST OF FIGURES

Contact addresses

wilken.hemme@hm.edu; benjamin.kormann@hm.edu

CC BY 4.0

References

[1] Congressional Budget Office. Usage Patterns and Costs of Unmanned Aerial Systems. Cbo study.

[2] Björn Müller. Die drohnenarmee Armee, 10 2024. https://www.reservistenverband.de/loyal/die-drohnenarmee/.

[3] Thomas Doll. Drohnen und Unbemannte Systeme: Gamechanger zur „Kriegstüchtigkeit“? *Hardthöhenkurier*, 40(5):75, 2024. ISSN: 0933-3355.

[4] Daniel Michaels. Ukraine’s drone attack exposes achilles’ heel of military superpowers. *The Wall Street Journal*, June 2025.

[5] Veronika Samborska. Scaling up: how increasing inputs has made artificial intelligence more capable. *Our World in Data*, 2025. https://ourworldindata.org/scaling-up-ai.

[6] Gordon E Moore et al. Cramming more components onto integrated circuits, 1965.

[7] Nobel Prize Outreach AB. The nobel prize in chemistry 2019, 2019. Awarded jointly to John B. Goodenough, M. Stanley Whittingham and Akira Yoshino "for the development of lithium-ion batteries". https://www.nobelprize.org/prizes/chemistry/2019/summary/.

[8] Silvus Technologies. Leading the mimo revolution, 2025. Company homepage describing Mobile-Networked MIMO (MN-MIMO) wireless communications systems for tactical and mission-critical applications. https://silvustechnologies.com/.

[9] Wilken Hemme, Jens Halbig, and Benjamin Kormann. Oben im dienst der natur: Untersuchung von UAV-sensoren zur wildtiererkennung und für den umweltschutz. In *Deutscher Luft- und Raumfahrtkongress 2024, Hamburg*, Bonn, 2024. Deutsche Gesellschaft für Luft- und Raumfahrt - Lilienthal-Oberth e.V. DOI: 10.25967/630459.

[10] Wilken Hemme, Benjamin Kormann, and Jens Halbig. Uav-based sensor comparison for fawn detection: Advancing wildlife conservation techniques. In *To be determined (e.g., proceedings of a conference or workshop)*. Faculty 04, Hochschule München University of Applied Sciences, Munich, Germany, 2024.

[11] Werner J. Marti. Fiber-optic drones are changing the war in ukraine. *Neue Zürcher Zeitung*, March 2025.

[12] Herbert A. Simon. A behavioural model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118. ISSN: 0033-5533. DOI: 10.2307/1884852.

[13] Kevin Dill and Dave Mark. Improving ai decision modeling through utility theory. Game Developers Conference, 3 2010.

[14] Jeff Orkin. Three States and a Plan: The A.I. of F.E.A.R. In *Proceedings of the Game Developers Conference (GDC)*. International Game Developers Association (IGDA), 2006. AI Summit.

[15] Van Parunak, Michael Purcell, and Robert O’Connell. Digital pheromones for autonomous coordination of swarming uav’s. 05 2002. ISBN: 978-1-62410-106-9. DOI: 10.2514/6.2002-3446.

[16] M. Helmert. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26:191–246, July 2006. ISSN: 1076-9757. DOI: 10.1613/jair.1705.

CC BY 4.0