Embedded reasoning for UAV operations: towards real-time efficiency and trustworthy autonomy

Floriano Scioscia*, Giuseppe Loseto†, Arnaldo Tomasino*, Ivano Bilenchi*, Filippo Gramegna*,
Saverio Ieva*, Agnese Pinto*, Eugenio Di Sciascio*, Michele Ruta*
* Politecnico di Bari – Via E. Orabona 4, Bari (I-70125), Italy
{nome.cognome}@poliba.it
† Università LUM "Giuseppe Degennaro" – Casamassima BA (I-70010), Italy
loseto@lum.it

***Abstract*—Unmanned Aerial Vehicles are increasingly used in a variety of applications. Operations control currently relies on ground control stations for data analysis and decision-making, incurring in latency issues and requiring expensive computing resources. This paper proposes the integration of Knowledge Representation and Reasoning techniques into UAV autopilot platforms to enable autonomous context management, improved operational efficiency, and real-time risk detection. Computational efficiency and explainability characterize the adopted approach, promoting trustworthiness of autonomous UAV units, as shown by two case study reports.**

***Index Terms*—Semantic Web of EveryThing, eXplainable Artificial Intelligence, Unmanned Aerial Vehicles**

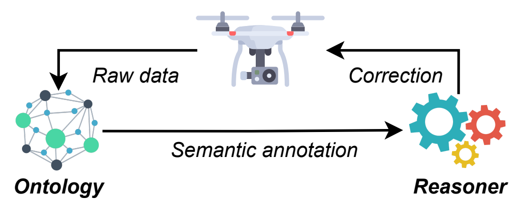

Fig. 1. Embedded reasoning for autonomous UAV operations

I. Introduction

Unmanned Aerial Vehicles (UAVs), a.k.a. *drones*, have been witnessing a significant expansion of their application domains, ranging from search and rescue operations to precision farming, hazard detection, logistics and more [1]. This progress has been facilitated by advancements in the miniaturization and integration of embedded sensors, microcontrollers, and programmable processing units. Sensing and decision-making capabilities of UAVs based on Artificial Intelligence (AI) leverage data analysis and inference procedures implemented in ground control stations, which communicate with the UAVs through wireless links [1]. Unfortunately, this approach suffers from latency issues, which can be incompatible with real-time use cases. Additionally, conventional Machine Learning (ML) methods often require dedicated onboard computing resources in order to run locally, leading to increased costs, weight, and energy demands. Furthermore, the poor transparency and explainability of several state-of-the-art ML methods impairs scenarios that require trustworthiness and accountability [2].

Knowledge Representation and Reasoning (KRR) techniques may be leveraged to address these challenges. They provide inherent interpretability by employing explicit machine- and human-understandable information modeling [3]. Nevertheless, most existing KRR tools have been designed for desktop-class devices, making them unsuitable for resource-constrained UAV embedded platforms. In this context, the emerging paradigm of the Semantic Web of Everything (SWoE) [4] aims to promote the integration of Semantic Web languages and KRR technologies in pervasive contexts. This involves designing KRR tools that are efficient and practical enough to be deployed on embedded microcontrollers.

This work highlights early results concerning the feasibility and benefits of embedding KRR capabilities into UAV autopilot platforms. By annotating sensor data and internal information using an ontology model, reasoning techniques can be leveraged for autonomous decision making. Most importantly, inference outcomes are associated with formal logic-based explanations, allowing UAV systems to justify their decisions [2]. The remainder of the paper is as follows: Section II describes the approach and, before conclusion, in Section III two illustrative examples showcase the usefulness of the proposal. They are taken from case studies implemented on the *Pixhawk* (https://pixhawk.org/) open standard platform for UAV autopilots with the *Apache NuttX* (https://nuttx.apache.org/) real-time operating system.

II. Proposed Framework

In UAV systems, achieving situational awareness and self-adaptation is crucial for effective operations. This entails monitoring the internal state of the UAV (such as kinematics, processing resource usage, and previous results) as well as contextual features derived from sensors and the environment (such as wind speed/direction, lighting conditions, and camera settings). By employing on-board reasoning capabilities, UAVs can automatically adjust their operations by determining the proximity of the current situation to critical states, and dynamically adapt to avoid them.

To support this reasoning process, sketched in Figure 1, information modeling plays a vital role: a reference Knowledge

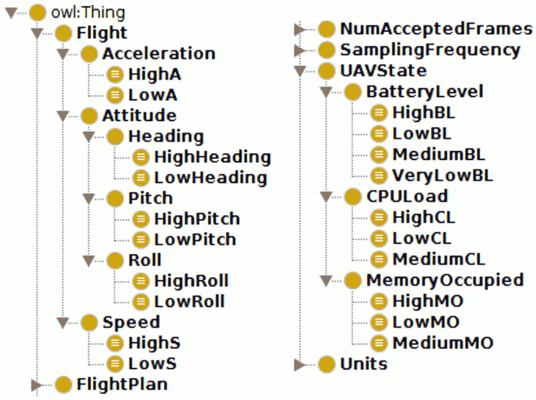

Fig. 2. Excerpt od the context awareness knowledge base

Base (KB) is created, including an ontology for characterizing (both internal and external) features and a set of individuals representing critical conditions. During UAV operation, the current state is dynamically annotated as a new individual, and its description is used to detect matches with any critical state by means of non-standard inference services for *semantic matchmaking* [3]. In case of a match, the UAV is prompted to proactively avoid the occurrence of the unwanted scenario. The whole process takes place locally, using the Cowl library [5] for KB management and the SWoE-oriented *Tiny-ME* inference engine [4]. Both tools have been explicitly designed to run on resource-constrained platforms and provide performance in line with embedded, real-time use cases.

III. Illustrative Examples

A. Context awareness

The UAV is used for crowd detection and avoidance via frame-based image analysis from nadiral camera payload. A comprehensive ontology has been developed, comprising classes that describe UAV parameters, and individuals describing critical scenarios that prevent getting reliable image frame acquisition or timely processing. An excerpt of the ontology is shown in Figure 2. To determine if the current conditions align with any of the critical scenarios, the current state of the UAV, e.g.:

HighRoll **and** *LowPitch* **and** *LowA* **and** *LowS* **and** *LowHeading* **and** *HighBL* **and** *HighCL* **and** *HighMO*

---

is compared with each critical scenario available in the Knowledge Base via semantic matchmaking, which returns a ranking of scenario profiles based on semantic similarity. In this particular case, the following profile passes the semantic similarity threshold, due to the currently high CPU load and memory occupancy:

HighCL **and** *HighMO*

To mitigate the impact of the identified critical scenario, a strategic decision is made to skip the processing of the next frame, alleviating the UAV's processing load and enhancing its overall operational efficiency.

B. Environmental hazard detection

This example focuses on the real-time identification of explosion risk levels using a UAV system, based on the European Union Directive 2014/34 (https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32014L0034). The ontology models crucial aspects of the UAV, the environment and the substances of concern, such as on-board sensors, actuators, detectable substances, dangerous levels, and critical atmospheric conditions specific to these substances. A periodic task continuously monitors the environment as follows: values are collected from on-board sensors; a concept expression representing the current scenario is constructed; semantic matchmaking is performed to compare the newly created individual with the risk profiles stored in the KB. Consider the following scenario and available risk profiles:

Scenario $\equiv$ *MediumConcentrationMethane* **and** *HighOxygenConcentrationMethane* **and** *LowVentilationMethane*
ExplosiveMethane $\equiv$ *HighConcentrationMethane* **and** *HighOxygenConcentrationMethane* **and** *LowVentilationMethane*
FlammableMethane $\equiv$ *MediumConcentrationMethane* **and** *HighOxygenConcentrationMethane* **and** *LowVentilationMethane*

The system identifies *FlammableMethane* as the nearest individual, indicating the presence of a flammable atmospheric condition, and a methane-related alert is raised. Human operators can immediately understand why the alert was raised, as allowed by the semantic-based concept expressions.

IV. Conclusion

This paper has proposed the integration of SWoE-oriented KRR techniques and tools into UAV systems, in order to improve operational efficiency and decision explainability w.r.t. conventional ML methods. Illustrative examples demonstrate the potential benefits in context-awareness and hazard detection. Future work will focus on systematic performance evaluations on relevant UAV platforms, expanding the scale and scope of applications, and exploring knowledge-based information fusion within UAV swarms as well as with vehicular networks and smart city infrastructures.

Acknowledgments

This work has been partly supported by project TEBAKA (TErritorial BAsic Knowledge Acquisition), funded by the Italian Ministry of University and Research.

References

[1] H. Shakhatreh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, and M. Guizani, "Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges," *IEEE Access*, vol. 7, pp. 48 572–48 634, 2019.

[2] High-Level Expert Group on AI, "Ethics guidelines for trustworthy AI," European Commission, Brussels, Tech. Rep., April 2019.

[3] M. Ruta, F. Scioscia, G. Loseto, A. Pinto, and E. Di Sciascio, "Machine learning in the Internet of Things: A semantic-enhanced approach," *Semantic Web*, vol. 10, no. 1, pp. 183–204, 2019.

[4] M. Ruta, F. Scioscia, I. Bilenchi, F. Gramegna, G. Loseto, S. Ieva, and A. Pinto, "A multiplatform reasoning engine for the Semantic Web of Everything," *Journal of Web Semantics*, vol. 73, p. 100709, 2022.

[5] I. Bilenchi, F. Scioscia, and M. Ruta, "Cowl: A Lightweight OWL Library for the Semantic Web of Everything," in *Current Trends in Web Engineering. ICWE 2022*. Springer, 2023, pp. 100–112.