Immersive AI assistance during eVTOL multi-agent ATC traffic routing

Jean-François Delisle, Simon Riendeau

CAE Inc.

Montréal, Québec, Canada

jean-francois.delisle@cae.com, simon.riendeau@cae.com

Clodéric Mars, Sagar Kurandwad

AI-Redefined

Montréal, Québec, Canada

cloderic@ai-r.com, sagar@ai-r.com

ABSTRACT

Advances in artificial intelligence (AI) such as natural language processing (NLP) and reinforcement learning (RL) are enabling breakthrough advancements in immersive flight training for both instructor-led and self-paced training, as exemplified by immersive Mixed Reality (MR)/Virtual Reality (VR) devices.

Introduction of MR/VR features will likely require dialog management and command and control capability in a self-paced training context. Adding an NLP-based cognitive agent acting as a virtual instructor and co-pilot provides the required immersion level to broaden the spectrum of self-paced training during the pilot's learning journey. Pilots receive instant feedback on their performance, explanations of the communication procedures, and progress tracking as they develop their skills.

Accent-tolerant advances in speech interaction are used to recognize radio transmissions from the ownship using NLP on a flight training knowledge base. Conversation agents increase student immersion and offer more realistic workloads by fully automating the air traffic control (ATC) function, freeing up the instructor to focus more on core observation and training performance, and allowing more automated flying without an instructor for MR/VR simulation training.

The complexity of real-world ATC communications, such as conditional clearances and instructions that include "give way," can only be simulated when the ownship is fully embedded with other traffic. An AI ATC module leverages a collaborative multi-agent framework to manage air traffic during real-time MR/VR flight simulation in a synthetic urban environment scenario.

This paper will explore issues around robustness in reinforcement learning and evolutionary optimization problems, alongside new results in collaborative multi-agent systems. These outputs will provide further results in nonlinear function approximation (e.g., deep neural networks) and optimization methods in stochastic environments. The paper will also study human factors during immersive MR training of emergency scenarios with an AI agent integrated into an electric Vertical Takeoff and Landing aircraft (eVTOL) flight training simulation and operation platform.

ABOUT THE AUTHORS

Jean-François Delisle, CAE Inc., AI Innovation Lead, Advanced Technology and Innovation
Jean-François has over 25 years of experience in software engineering, AI, and data solutions. As a member of the Advanced Technology & Innovation department, he provides new solutions for the training and learning ecosystem. His research focuses on AI solutions that help optimize and adapt training delivery and the learning experience through human behavior and performance analysis. Jean-François joined CAE in 2010 and his current mandate is to define flight training technology strategies and disruptive innovation using data analysis and AI capabilities in Advanced Air Mobility and eVTOL engineering solutions. He earned a PhD in artificial intelligence and cognitive science for adaptive flight training under the supervision of Prof. Andrea Lodi, Canada Excellence Research Chair in Data Science for Real-Time Decision Making at Polytechnique Montréal.

Clodéric Mars, AI-Redefined, VP of engineering
Clodéric has been building AIs since 2006 with one overarching goal: fostering collaboration between humans and AIs. At National Institute for Research in Digital Science and Technology INRIA, then at Golaem and MASA Group, he worked on explicit AI techniques applied to video games, simulation and special effects. In 2015, he co-founded Craft AI to focus on making machine learning (ML) explainable for time series prediction and analysis. In 2020, he joined AI Redefined to accelerate the development of the Cogment orchestration platform to further the synergy between humans and AIs. Over the years, Clodéric contributed to the development of AIs through collaborations with artists to populate movie shots, designers to understand transportation infrastructure usage, instructors to create realistic and purposeful military training, and engineers to operate energy infrastructure.

Sagar Kurandwad, AI Redefined, Researcher and Developer
Sagar is passionate about creating novel AI algorithms. He previously served as a lead researcher at Dailyhunt/VerSe Innovation, a leader in ML and deep learning technologies for content customization. Sagar began his career as a data scientist at Personagraph, Housing.com, and Tata Consultancy Services, and an instructor for ML and data analytics courses at BrainStation.

Simon Riendeau, CAE Inc., AI Strategy Data Science
Simon has been building software since the age of 11, and professionally since 2000. With a background in engineering and experience in aircraft systems and avionics simulations at CAE as well as video game development at Ubisoft Montréal, his current focus is on innovative practical applications of AI from data analytics to virtual assistants. Simon holds a bachelor's degree in mechanical engineering (Aeronautics) from Polytechnique Montréal and a graduate diploma in engineering management from Sherbrooke University.

Immersive AI assistance during eVTOL multi-agent ATC traffic routing

Jean-François Delisle, Simon Riendeau

CAE Inc.

Montréal, Québec, Canada

jean-francois.delisle@cae.com, simon.riendeau@cae.com

Clodéric Mars, Sagar Kurandwad

AI-Redefined

Montréal, Québec, Canada

cloderic@ai-r.com, sagar@ai-r.com

INTRODUCTION

Electric Vertical Takeoff and Landing vehicles (eVTOLs) have real potential to shape the future of Advanced Air Mobility (AAM), one in which these small and non-polluting aircraft revolutionize Urban Air Mobility (UAM) and Regional Air Mobility (RAM) transportation. While the future may see eVTOLs controlled from a distance through a Remotely Piloted Aircraft System (RPAS) (De Lellis et al., 2019) or Uncrewed Aircraft System (UAS) (Jeffrey Homola, 2016), on-board single pilot operation is expected to be the norm for several years. According to (Uri Pelli, 2020), AAM will increase existing pilot demand up to an additional 60,000 by 2028. Fast-tracking the training of the future pilots needed to operate eVTOLs is crucial, as is putting in place the cockpit automation required to limit workload while ensuring the safety of autonomous flight.

According to the EASA Artificial Intelligence Roadmap 2.0.(EASA, 2023a), industrial actors expect the first crew assistance/augmentation developments over a 2022-2025 time frame. Automation will ramp up to single pilot operations through human-machine collaboration from 2025 to 2030, culminating in remotely operated flights with human supervision or fully autonomous flights foreseen arriving after 2035.

NASA (Charles T. Howell, 2019) and the Federal Aviation Administration (FAA) are working together to perform research on UAM technologies. Under the Flight Demonstration Capabilities (FDC) Project, which supports the use of Autonomous Flight Safety Systems (AFSS), UAM Testbed Flight Research Aircraft will be used in the development and testing of automation, sense and avoid, air traffic control (ATC), aeronautics, energy storage and other related technologies. The UAM Testbed Research Aircraft, a Cessna LC40, is equipped with computers, sensors and software to serve during planned UAM research projects. Experiments on air traffic sense and avoid, automatic ground collision avoidance systems and Automatic Voice Recognition and Response System are planned to be executed. The goal of this last experiment is to automatically recognize ATC voice commands and generate automatic aircraft response.

Security, Safety and Regulation in Urban Air Mobility

UAM introduces new types of aircraft and operations that require appropriate regulations and certification processes and the implementation of security and safety measures. eVTOLs providing air taxi service must meet strict safety standards involving rigorous certification processes. Ensuring the safety and reliability of UAM operations is essential for acceptance by regulators and the public. UAM operations have the potential to introduce additional noise and environmental challenges that require the development of quieter aircraft, the optimization of flight paths and the transition to sustainable energy sources.

For this purpose, (Md. Shirajum Munir, 2023) propose a cyber-physical safety analyzer framework driven by artificial intelligence (AI). The framework uses AI algorithms, including a decision tree, random forests, logistic regression, K-nearest neighbors (KNN) and long short-term memory (LSTM), to predict and detect cyber jamming and spoofing attacks through a Received Signal Strength Indicator (RSSI), remote RSSI and voltage fluctuations. The developed framework of the AI algorithm then analyzes the conditional dependencies based on the Pearson's correlation coefficient identify the cause of potential attacks based . The experiment results show the efficacy of the proposed framework at around 99.9% accuracy for jamming and spoofing detection using a decision tree, random forests and KNN, while efficiently finding the root cause of the attack.

Urban Air Mobility Operations

UAM faces several infrastructure challenges for air travel in urban areas, including the establishment of intelligent air traffic management (ATM) systems, vertiports and charging stations, and their integration with existing transport networks. Traditional ATC systems are not designed to handle the density and dynamics of UAM operations. With effective management of air traffic in urban areas a major challenge, the development of robust UAM ATM systems is crucial, they have a large benefit of leveraging machine learning for incident risk predictions (Xiaoge Zhang, 2019). The paper of (Brock Lascara, 2019) explores these challenges and presents operational concepts for integrating highly automated UAM operations into the FAA's National Airspace System. The proposed framework can serve as a starting point for concept evaluations of an Airspace Integration Framework and is articulated around Augmented Visual Flight Rules, Dynamic Delegated Corridors, Automated Decision Support Services and Performance-Based Operations. Their research questions include impacts to air traffic managers controlling that airspace; impacts to other visual flight rules traffic in that airspace and other instrument flight rules traffic in the vicinity; decision support capabilities for ATC, UTM and operators; and procedural changes/additions needed to enable operations.

The NASA Operational Concept of (George Price, 2020) describes a community vision of the projected evolution of urban air service, with vehicles capable of carrying one or more passengers. The paper identified the challenges and gaps that must be addressed to enable the UAM vision. To categorize these challenges and their dependencies, the authors developed a framework comprising five pillars: Vehicle Development & Production, Individual Vehicle Management & Operations, Airspace System Design & Implementation, Airspace & Fleet Operations Management and Community Integration. The OpsCon envisions eVTOL aircraft in operation, from vertiports to ATM concepts and collision avoidance systems, that assure safe and efficient operation within the airspace.

Air Traffic Management in the Context of Urban Air Mobility

Decision support systems for UAM operations will need to leverage and exchange information with existing and conventional ATC services. A high level of automation support has been increasingly adopted to support tactical deconfliction tasks in UAM (Shulu Chen, 2023). ATM systems must be highly automated and intelligent. A highly automated Urban traffic management (UTM) system based on AI will need to address issues related to the explainability of intelligent algorithms in environments where human operators are involved in safety-critical decisions.

The research paper by (Yibing Xie, 2020) investigates the adoption of ATM AI for incidents and accident risk prediction through the XGBoost algorithm. The study focuses on explaining the trained AI model and the predicted results. Moreover, considerations are made on the most promising strategies to strengthen the trust between the ATC and the system through the redesign of the interface of Human-Machine Interaction. With the increasing role of AI, the complexity of machine learning (ML) "black boxes" has been increasing, which then raises the need for greater transparency through an Explainable AI (XAI). The research introduces the aviation incident and accident prediction model adopting the XGBoost algorithm, which is part of the ATM/UTM Decision Support System (DSS). The research aims to improve the development of an ATM/UTM application by using AI algorithms, exploring the explanation methods of the results given by the model with Explainable AI (XAI) and using SHapley Additive exPlanations as the post-explanation method to achieve the trustworthiness and reliability of the AI system. Ensemble learning including boosting (XGBoost), bagging and stacking are used to build predictive and explanatory models using the XAI method for model transparency and post-hoc explainability.

UAM Flight Training and Intelligent Human-Machine Interface

According to CAE research (CAE, 2021), UAM will create additional demand for pilots, who will need additional skill sets to operate in the new challenging UAM traffic environment. Modern learning techniques include adaptive learning with self-paced student-centric training programs using immersive training technologies. Augmented Reality (AR), Virtual Reality (VR) and Mixed Reality (MR) are incorporated into the initial training program to ramp up pilot qualification to meet demand while maintaining a low carbon footprint and improving training cost efficiency.

In UAM applications, additional physiological sensors can be leveraged to improve Human-Machine Interface (HMI) efficiency with technology such as voice recognition and biometric sensors. These sensors are used to analyze cockpit communications, track visual scan patterns and measure the pilot's cognitive load. Machine-based monitoring/enhancement helps prevent cognitive overload when increasing the level of autonomy in decision support systems to avoid problems, such as trust and loss of situational awareness in HMIs.

Research Objectives

This paper will study the efficiency of immersive MR training during emergency scenarios with an AI agent integrated into an eVTOL flight training simulation. The eVTOL Cognitive Agent is an AI-powered conversation agent that can help a pilot train on an eVTOL aircraft simulator without the need for a dedicated instructor. We aim to provide AI performance metrics on which to evaluate a Natural Language Understanding (NLU) Dialog Manager with a real-time gateway to flight simulations.

Intelligent agent dialog from the MR flight simulation will be integrated into collaborative multi-agent reinforcement learning for AI ATC traffic routing during eVTOL emergency landing scenarios. The paper also explores issues around robustness in reinforcement learning and evolutionary optimization problems, alongside new results in collaborative multi-agent systems. These outputs will provide further results in nonlinear function approximation (e.g., deep neural networks) and optimization methods in stochastic environments. The objectives are to:

- Enable student pilots to interact with an urban environment populated by other eVTOL aircraft that depict realistic behaviors and interactions.
- Implement an intelligent ATC entity with which both human pilot and autonomous aircraft interact to request clearances and routing.
- Showcase these systems in an emergency landing training scenario requiring unplanned synchronization between the human pilot, ATC and the autonomous aircraft.

The simulation training models a multi-agent system, including the Air Traffic Controller and eVTOL pilots (autonomous or human). This multi-agent system was implemented with the following objectives:

- Ability to train using reinforcement learning with the intent of injecting human demonstration or feedback for fine-tuning.
- Ability to operate in a simulation suitable for pilot training.
- Explore issues such as robustness in reinforcement learning and evolutionary optimization problems, alongside new results in collaborative multi-agent systems.

Additional objectives were to begin work toward complying with the recently published EASA guidelines. The EASA Concept Paper for Machine Learning (EASA, 2023b) for Level 1 Artificial Intelligence (assisting humans) and Level 2 Artificial Intelligence (human-machine collaboration) aims to proactively address forthcoming EASA guidelines and safety standards pertaining to ML applications with safety implications. The paper provides guidance to applicants who are incorporating AI/ML technologies into systems designed for safety or environmental purposes.

The EASA guidelines cover the following building blocks that lead to Trustworthy AI:

- AI Trustworthiness Analysis
- AI Assurance
- Human Factors for AI
- AI Safety Risk Mitigation

METHODOLOGY

The Mixed Reality Flight Training Device

The MR training device is an affordable, visually immersive flight simulation designed to deliver a realistic and interactive virtual environment. The device offers hardware flight controls and a simplified vehicle operation cockpit representative of eVTOL aircraft. The device features a high-fidelity optical head-mounted display, model JVC S1, as shown in Figure 1. The Generic eVTOL Simulation runs on a virtualized on-premise computing environment with immersive systems such as: Seat Vibration, 3D Sound, and Short-Stroke Motion. The cloud computing platform is used for data and ML engineering. The flight simulation is monitored and controlled by a remote Instructor Operator Station (IOS), while flight telemetry is collected for simulation playback and analytics.

Figure 1 – JVC S1 high-fidelity optical head-mounted display on a motion device with flight control and cockpit replica

Experimentation Design

The eVTOL Cognitive Assistant powered by AI is parametrized to support conversation that can help pilot training on an eVTOL aircraft simulator, as shown by Figure 2. The ATC scenario is deployed in a synthetic environment. The Open Geospatial Consortium common database (OGC CDB) interacts with Computer Generated Forces (CGF) systems that can interface with integrated, human-in-the-loop simulators interoperable with different open standard High-Level Architecture communication protocols in federated exercises between multiple entities. Collaborative multi-agent reinforcement learning, interacting with a gateway to the simulation, executes AI Air Traffic Controller traffic routing during eVTOL emergency landing scenarios with intelligent agent dialog. Figures 2 and 3 illustrate the major components of the systems and their interconnections with each other and with the simulation services, using a mix of proprietary and industry-standard interchange protocols.

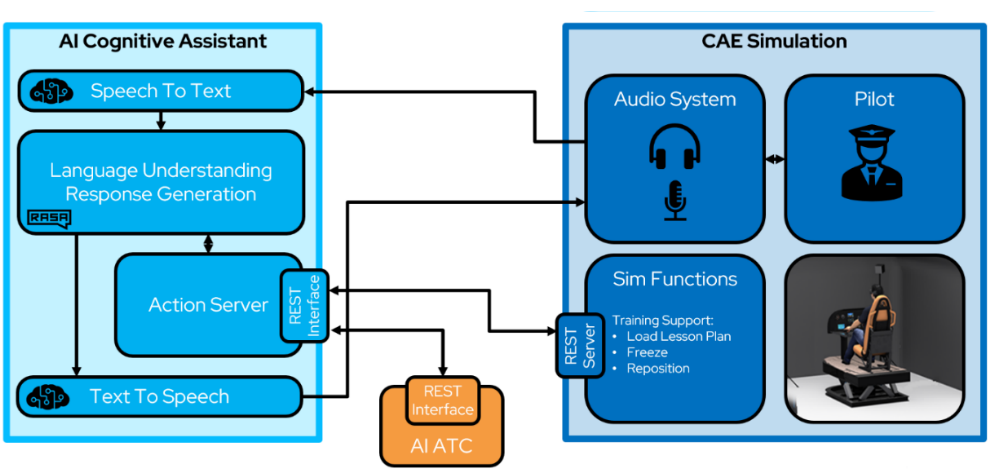

Figure 2 – Cognitive Assistant High-Level Architecture

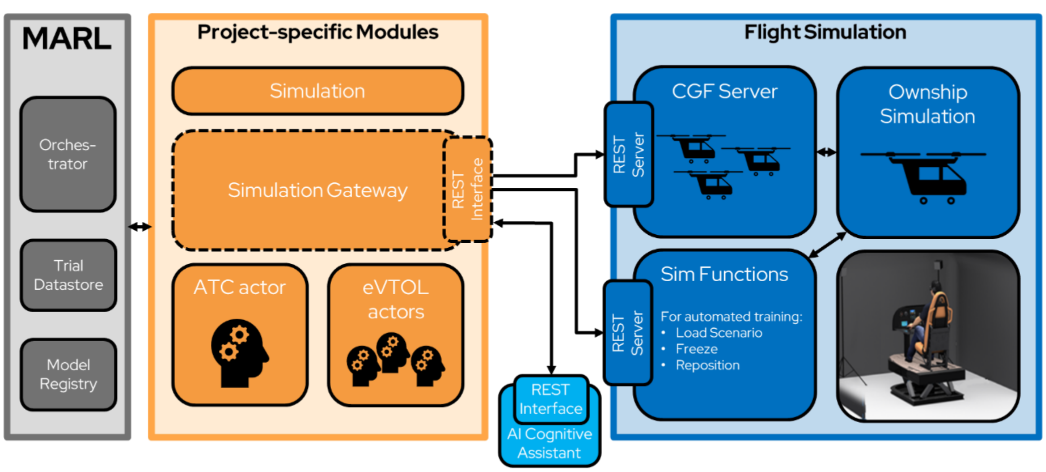

Figure 3 – Simulated ATC High-Level Architecture

The Cognitive Agent

The Cognitive Agent serves as the front end to the student, as the entity with which the student can converse and utter requests or ask questions and receive answers. In practice, the Agent's capabilities can be attributed to these distinct roles: Virtual Instructor, Virtual Co-pilot and Simulated Air Traffic Controller. For each role, a dedicated logic module uses various levels of sophistication, from expert systems to carefully orchestrated AI models, to interface with the front-end layer that supplies the voices. All is coordinated to provide the student a cohesive conversation experience.

Voices

While each role serves a purpose, it made sense to combine the roles of Virtual Instructor and Virtual Co-pilot into a single persona with a given voice, identified by the call sign *Cassie*. The Simulated Air Traffic Controller, a role located at a virtual fixed location outside the aircraft, was assigned a persona and voice with the informal call sign, *Bob*. The creation of these two distinct personas/voices helps the user distinguish with whom they are interacting.

Role: Virtual Instructor

The Cognitive Agent takes on tasks an instructor would typically perform, which can be further split into two roles: Virtual Coach and Flight Training Assistant. The Virtual Coach supplies expert guidance and insight to the student, including providing live feedback and session briefing/debriefing. The Flight Training Assistant can interact with the simulator itself to control the training session by interfacing with the remote IOS to:

- Load lesson plans.
- Toggle the simulation freeze.
- Reposition the simulator to a requested location.

Role: Virtual Co-pilot

As a stepping-stone between two-pilot and single-pilot operation, the Cognitive Agent has the ability to act as a Virtual Co-pilot. This role interfaces directly with the aircraft system to further automate cockpit operations and reduce the load on the pilot. Capabilities are diverse and can include:

- Operating aircraft functions (control the landing gear, air conditioning, light, etc.).
- Running through checklists.
- Monitoring flight operations.
- Assisting with navigation.
- Communicating with the ATC on behalf of the pilot.

Role: Simulated Air Traffic Controller

The Cognitive Agent also provides the voice of the Simulated Air Traffic Controller, and the conversation system actually supports discussion between the Virtual Assistant and Simulated Air Traffic Controller personas.
Role capabilities focus on managing traffic within a simulated area in complex scenarios, including:

- Monitoring air traffic.
- Dispatching and rerouting aircraft.
- Managing emergency situations.
- Communicating with the student and with the other simulated pilots, for realistic radio chatter.

Artificial Intelligence

There are two AI/ML components in the eVTOL cognitive agent application: AI Speech service and Rasa Open Source. AI Speech service provides cloud-based, ready-to-use speech-to-text and text-to-speech capabilities where minimum model fine-tuning is required. The speech-to-text capability enables real-time transcription of audio streams into text. The baseline speech-to-text model has been trained with Microsoft-owned data, which can identify generic, day-to-day language. Our custom speech-to-text model is built upon this baseline with enhancements on ATC special terms and phrases. The text-to-speech capability converts text into humanlike, synthesized speech. The baseline text-to-speech model is used, and there is no model customization.

Natural Language Processing

The central component of the eVTOL Cognitive Assistant application is the Natural Language Processing (NLP) capability powered by the Rasa Open Source framework. Rasa is an AI platform that provides the building blocks to create the virtual assistant, with modules performing the NLU and connecting to other services, such as the simulation services presented by the architecture of Figure 3.

The NLU processing of Rasa is customized via training data in the form of text files that define *intents* and *stories*. Each of the possible commands and questions the chatbot can recognize is an *intent*, whereas *stories* are used to define the possible flow of a conversation and the applicability of the intents at any given moment. Multiple versions of utterances are supplied for each intent, with the possibility to include parameters that will be isolated. For instance, both "reposition to ten thousand feet" and "please repos to ten k" will be classified as a "reposition" intent with a parameter of "10,000 ft" for use in the subsequent logic.

Thus, each message incoming, as text, is processed by a sequence of components. It is analyzed in the context of the stories and mapped to a valid intent, if any. Additional custom modules handle the appropriate response and/or interactions with other services. In the case of a request to reposition, the system may reply with an acknowledgement, followed by a signal to request the simulation to perform the reposition then an additional vocal notification when the process is completed.

Collaborative Multi-Agent Reinforcement Learning

The target domain and use case is modelled as a multi-agent system (MAS), as shown in Figure 4. Two types of agents were identified:

- Air Traffic Controllers, in charge of analyzing the airspace, and communicating with the eVTOL to dispatch commands, issue altitude clearances, and identify heading and routing patterns. This agent is also in charge of helping maintain aircraft separation.
- eVTOL pilots handling communication with the Air Traffic Controller and steering their eVTOL to follow the flight plan. Two versions of the eVTOL pilots were developed: autonomous and "piloted." The autonomous pilot manages the structured communication with the Air Traffic Controller and updates local waypoints to be followed by the simulated eVTOL entities. The "piloted" version replicates the human pilot in ownship and relies on the Cognitive Assistant to act as the translation interface in the structured communication with the Air Traffic Controller and the voice communication with the Virtual Co-pilot.

The target simulation presents a scenario with airport and urban vertiport environment interactions. In the initial scenario, an airport to urban vertiport route, a malfunction forces the pilot to interact with Air Traffic Controller for an emergency landing at vertiport. Aerial Traffic reacts to ownship and ATC dispatches command, behaving in a non-deterministic, scenario-style fashion.

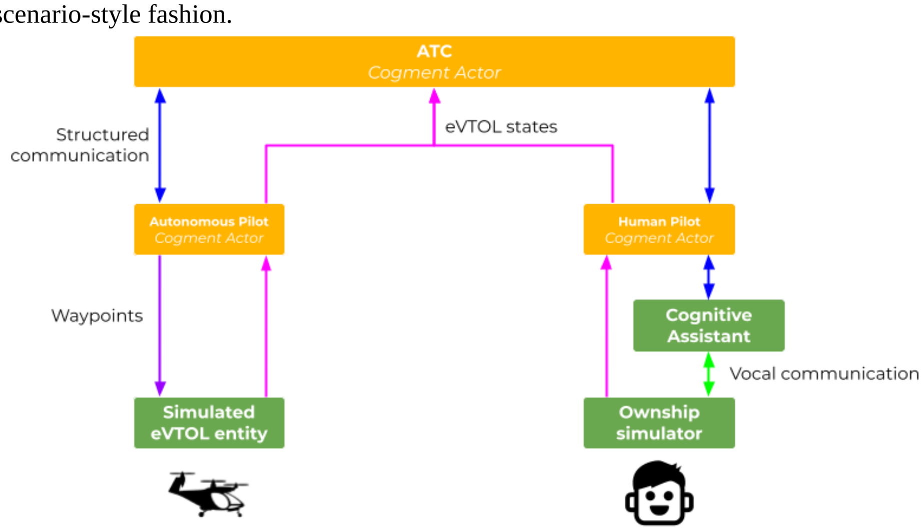

Figure 4 – Multi-agent modeling (orange identifies the agents implemented as Cogment actors, green identifies components belonging to the other systems) and component communication (blue for structured message communication, pink for airspace states, green and purple for waypoints)

To implement, train and operate this MAS, we used Cogment (Sai Krishna Gottipati, 2021), a platform designed to build, train and operate AI agents in simulated or real environments shared with humans. In particular, we explored how Reinforcement Learning (RL) can be leveraged to train this collaborative MAS. The results of these experiments are presented in the Result section.

Figure 4 illustrates and identifies the different components of the software architecture:

- Blue identifies the base Cogment components — the orchestrator, trial datastore and model registry.
  - The orchestrator is responsible for the execution of trials over actors and environments.
  - The trial datastore listens for the data generated by the trial, e.g., the trajectories of the different actors and environments. It stores this data and makes it available for AI training or other purposes online or offline.
  - The model registry hosts trained models and makes them available to actors.
- Orange identifies components implemented using the Cogment software development kits include the actors in charge of taking action from received observations; the environment that applies the actions and generates situational observations; and the runners, which execute on-demand trials for interactive operations of the platform (i.e., a pilot playing a scenario) and batch queries to train AI models. Those components communicate with the Cogment components using Cogment's gRPC API. We discuss the different actors envisioned RESULTS section.
- Green identifies the simulation gateway, a set of HTTP servers able to control the simulation, the Computer Generated Forces and the Cognitive Assistant as described in the Experimentation Design section.

Cogment relies on the formalism of actor classes derived from Markov Decision Processes. Each class is defined by observation and action spaces (i.e., what it perceives of the environment and what actions it can perform in the environment). Agent modelling is mapped into two Cogment actor classes:

- Aircraft pilot actors represent the decision-making of eVTOL aircraft. Their observation space includes other eVTOL aircraft in their perception range and messages sent to them by the Air Traffic Controller. Their action space includes local waypoints and communication with ATC (in particular, flight plan updates).
- Air Traffic Controller actors represents ATC decision-making. Their observation space includes the state of the airspace they are responsible for and communication with aircraft. Their action space includes communicating with the aircraft.

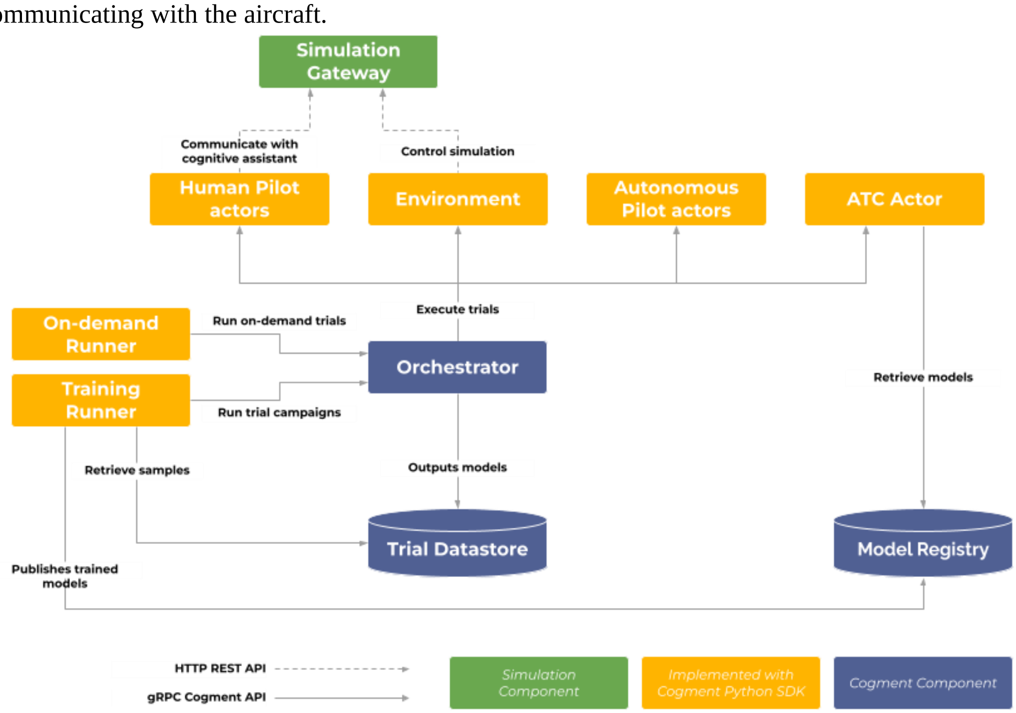

Figure 5 – Cogment-based micro-service architecture

Class defines what each actor can do; implementation defines how the actor performs the action. Cogment enables multiple implementations of each actor to coexist transparently. For the aircraft pilot actors, we have two implementations:

- The ownship, which is operated by human trainees interactively, does not generate waypoints, as the human pilots already have control over the aircraft outside of Cogment.
- A doctrine-based scripted behavior.

Trained implementations of the aircraft pilot actors, using Deep Reinforcement Learning (DRL), were considered but were not implemented as part of this work because sufficient perceived diversity, realism and reactivity were achieved by training the Air Traffic Controller, which is the focus of the following section.

For the Air Traffic Controller actor, we implemented a hybrid approach leveraging a RL-trained agent, observation preprocessing and heuristic implementation of high-level tasks. This approach is shown in Figure 5. The first part of the actor implementation addresses reception of inbound messages from the eVTOLs and the update of the airspace to the latest known state of each eVTOL, including:

- Current location.
- Final destination.
- Current route.
- Latest takeoff, landing or routing request.
- Incident declaration.

This list of eVTOL states was used as an observation for the RL-trained agent. The output action space includes a combination of identifying the eVTOL and the task to be performed: “Handle Incident,” “Handle Routing,” “Authorize Takeoff,” “Authorize Landing” or “No Action.” A downstream process was then applied. “Handle Incident” assigns the nearest vertiport for an emergency landing to the eVTOL and updates the final destination of the other eVTOLs navigating toward the same vertiport. The “Handle Routing” process computes a new route toward the final destination. Other processes were simply about sending the right message to the eVTOL pilot actor.

This architecture makes the role of the RL agent about prioritizing which action to take, while leaving the execution of the action to a heuristic process. Our RL agent training uses the Advantage Actor Critic (A2C) algorithm(Volodymyr Mnih, 2016) which we selected for its simplicity and stability. The A2C algorithm combines the advantages of both policy-based and value-based methods. In A2C, an actor network learns to choose actions based on observed states, while a critic network estimates the expected rewards or values of those states. The actor receives feedback from the critic in the form of advantages, which indicate how much better or worse the chosen action was compared to the expected value. This feedback helps the actor update policy, while the critic network updates value estimates based on the temporal difference error between predicted and actual rewards. By iteratively improving both the policy and value estimates, A2C enables more stable and efficient learning in RL tasks.

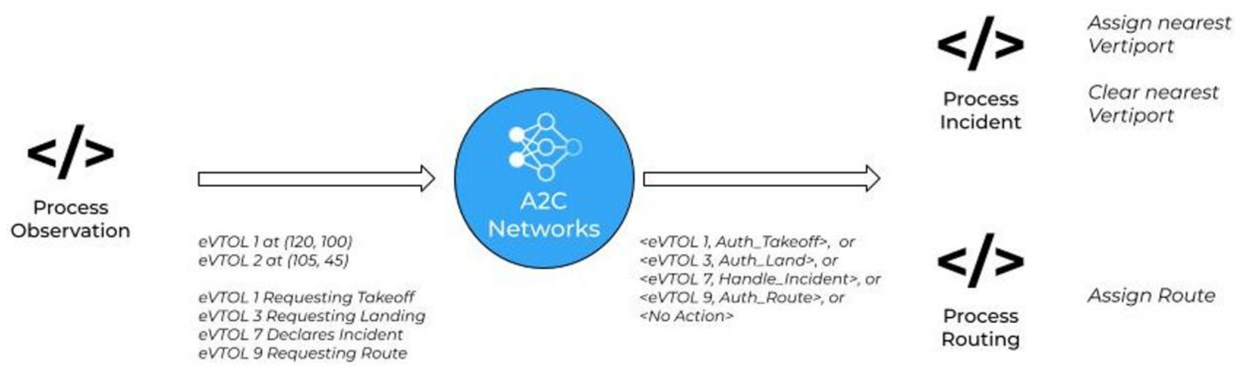

Figure 6 - Internal architecture of the Air Traffic Controller Actor leveraging A2C

We divided the A2C network into three smaller units. The first network processes features from the scene, such as the number of parking spots left in each vertiport, eVTOL routes, etc. The second network deals with the state of the communication channel, whereas the third process the eVTOLs locations, requests, intra-eVTOL distances among some other eVTOL states.

The overarching A2C network concatenates the processed vectors from each of the individual networks to be mapped to the action space. This network architecture is dependent on the number of considered eVTOLs, destination vertiports and route waypoints. However, these numbers should be considered as upper bound, as the agent is usable with fewer eVTOLs, vertiports or waypoints by zeroing the matching input values and ignoring the matching output values.

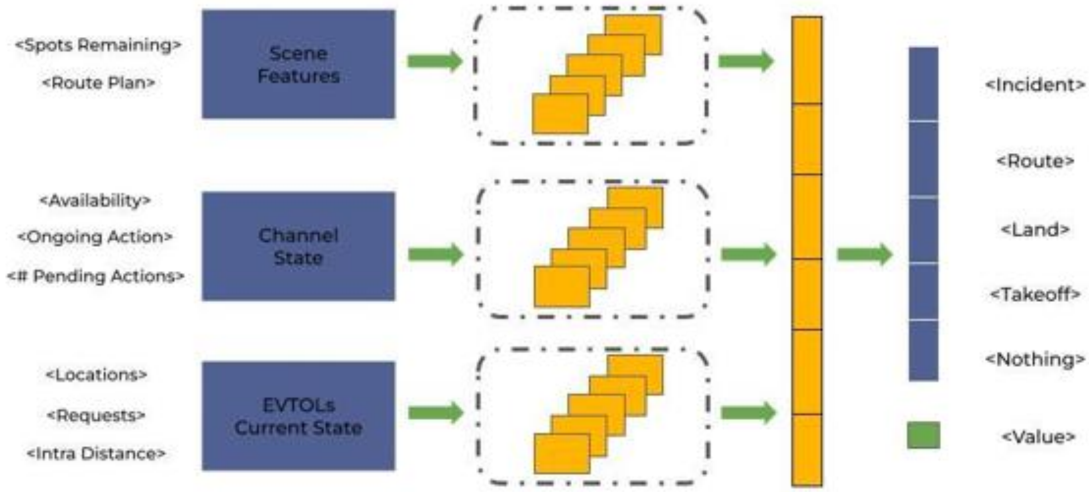

Figure 7 - ATC A2C network architecture

It is important to note that neither the network nor the pre- and post-process design or configuration consider whether the flight dynamics of specific aircraft or the environment. The only assumption is that all the aircraft types are equivalent, as the actor does not have any way of telling them apart. The Air Traffic Controller actor was trained by running a campaign of Cogment trials as a "Training Runner," as described in Figure 7. Each trial was configured with a randomized list of eVTOL having randomized destinations, with the trial ending when all eVTOLs reached their respective destinations. The reward function was designed to heavily penalize collision and separation break to encourage the safety of the eVTOL flights. Trial length also carried a penalty to encourage the Air Traffic Controller to help eVTOLs reach their destination and thus end the trial as quickly as possible:

Events	Reward/Penalty
Collision between eVTOLs	-2
Intra-eVTOL Distance <= 50 m	-1
Every Step	-0.05

Implementation of EASA AI Guidelines

The Virtual Instructor, Virtual Co-pilot and Simulated Air Traffic Controller were designed with safety in mind, the objective to facilitate self-paced training to achieve the goals of training pilots faster and better than before. This goal is accomplished by leveraging AI/ML technologies and EASA AI guidelines to provide a framework that ensures the technologies themselves are reliable and trustworthy. As such, many of these guidelines have been incorporated in the development pipeline of the solution. Examples include:

- Identifying the level of the AI application, its capabilities and its users.
- Ensuring that procedures avoid causing bias in the trained models.
- Assessing the risk of deskilling the users and mitigating it appropriately.
- Ensuring the correct data is collected, managed, processed, labeled, and used to train models.
- Documenting the solution.
- Ensuring appropriate configuration management of the solution.

RESULTS

Model training results are currently logged and recorded in the DevOps pipeline. We evaluate the NLU component using cross-validation. Cross-validation automatically creates multiple train/test splits and averages evaluation results on each train/test split. Intent recognition performance metrics and confusion matrix is shown in the figure below:

	intent recognition	entity extraction
Precision	0.989	: 0.956
Accuracy	0.989	: 0.968
F1-score	0.989	: 0.947

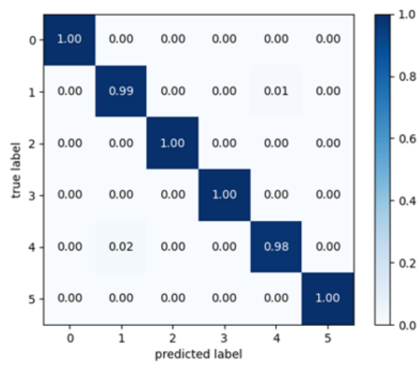

Figure 8 - Intent Recognition/Entity Extraction & Confusion Matrix

MARL Evaluation Metrics for ML Performance Monitoring

Experiments were conducted to evaluate the realism of the resulting air traffic behavior both qualitatively, in the real-time simulation environment, and quantitatively using domain metrics. The Air Traffic Controller was benchmarked against a greedy heuristic policy (first request-first service) baseline on a few scenarios with the number of eVTOLs, vertiports and waypoints varying between 1 and 10. The figure below shows the increase in the AI Air Traffic Controller's expected total rewards +/- 3 standard deviations with gains achieved through repeated experience across scenarios. One full training run lasted about 4 hours on a workstation with an Intel i9-9900K with 8 physical cores and 32 GB of RAM. No graphics processing unit was used for this experiment.

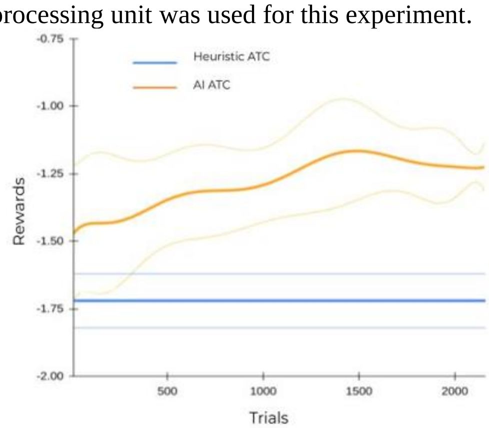

Figure 9 – Rewards evolution during the Air Traffic Controller actor training

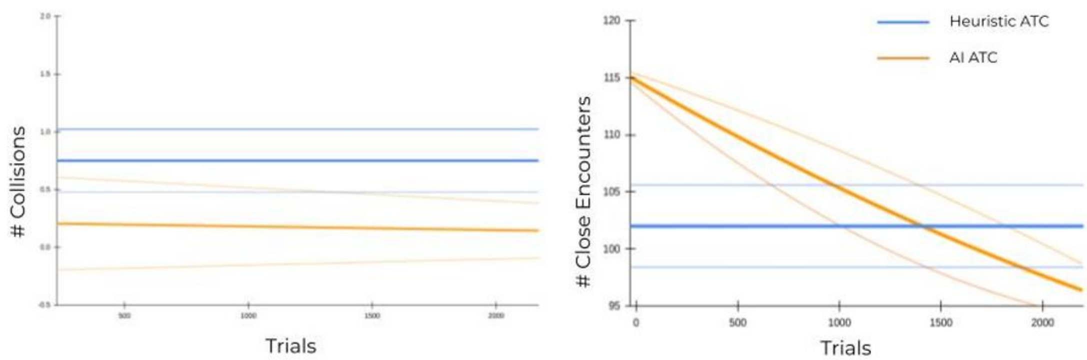

Figure 10 – Safety metrics evolution during the Air Traffic Controller actor training

The trained AI Air Traffic Controller's policy turns out to be a safer ride as compared to that of the heuristic counterpart. The number of collisions in a scenario directed by the AI Air Traffic Controller remains practically non-existent, with a drastic fall in the number of close encounters between eVTOLs (i.e., intra-eVTOL distance) as well. In comparison to the Heuristic Air Traffic Controller, the safety outcomes are considerably better.

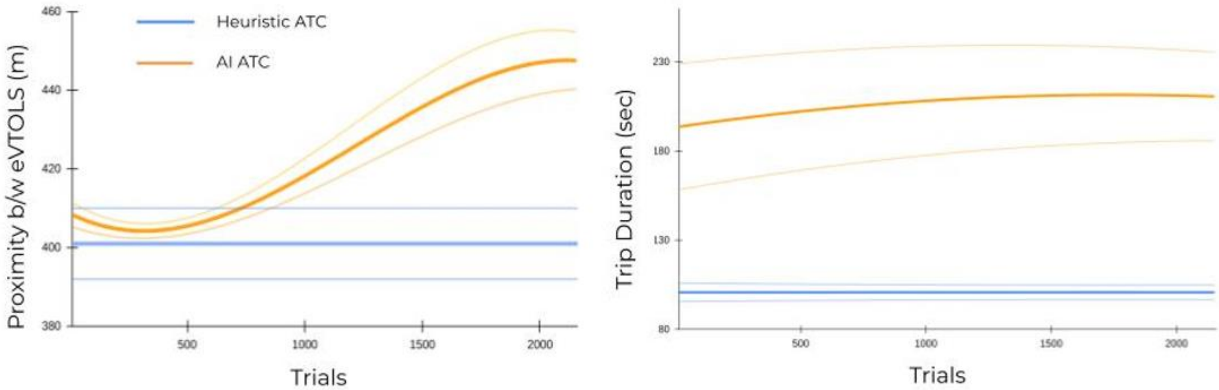

Figure 11 –AI Air Traffic Controller actor learns over time to prioritize safety

Since the trained AI Air Traffic Controller agent learns to maximize safety, proximity between eVTOLs is considerably greater during trials as compared to the eVTOLs directed by the Heuristic Air Traffic Controller. However, this causes the eVTOLs to take longer to complete their trips. These results show that the AI Air Traffic Controller was able to consider the safety objectives expressed as part of the reward function while fulfilling the assigned mission. Qualitatively, we could observe that the main difference between the greedy heuristic behavior and the trained one is that the latter tends to delay takeoff requests. This means the AI Air Traffic Controller was able to learn to anticipate conditions causing closing close encounters, in particular the traffic density around vertiports, and adapt behavior accordingly. We implemented the behavior of this multi-agent system using a hybrid approach that combined carefully crafted, but ultimately simple, heuristic processes involving the eVTOL pilot actors and the pre- and post-processing of the AI Air Traffic Controller actor. Conversely, we limited DRL for task priority. We demonstrated that, by encoding safety rules in the reward function, we could obtain high-fidelity behavior. Other desirable aspects could be considered by combining relevant metrics with the reward function. Furthermore, this approach is adaptable to new scenarios and, more importantly in the early day of the eVTOL industry, new aircraft dynamics through the retraining of the policy.

Capability Analysis on the Implementation of EASA AI/ML Guidelines

Engineering processes are now equipped with human factor evidence in a framework aligned with ethical standards. We discovered that the W-model (EASA, 2023b) of learning quality assurance allows traceability both in the engineering and training phases. This gives us the ability to analyze the human factors involved in the pilot's learning process with a higher level of trust in a human-machine-teaming context. The framework offers us a better ability to analyze immersive and conversational person-machine interfaces in the operation of eVTOL in an urban environment.

CONCLUSION

The objective of AAM activity is to transport passengers efficiently and safely in an urban environment, offering new challenges to pilots and traffic management systems. There is a need to develop technology to create a feasible, operable AAM transport system for autonomous or semi-autonomous vehicles using AI. MR flight training devices with artificial capability provide high-fidelity simulations to train individuals with the appropriate level of immersion and realism of a complex scene involving air traffic management operations in a synthetic urban environment. This can be leveraged to ramp up pilot qualification to meet demand, contributing a low carbon footprint and improving training cost efficiency.

As future work, we aim to address challenges in behavior realism of the Cognitive Agent by enhancing the natural behavior features. Speech should align speed, volume, tone, and formulation with the pilot mindset and the context. The level of realism and quality of results will depend on the availability of source data. We suggest procuring authorized video of exemplary co-pilot and coaching behaviors, as well as written requirements/expectations from subject matter experts. Emotional intelligence behavior can modify the Agent's own speech and intervention based on various criteria, such as the pilot’s state of mind and criticality of the situation, which are factors in the ability to accurately replicate stress normally expressed by a co-pilot in an emergency scenario. We also plan to consider the capability of making mistakes and the capability to configure the stress level set by default, scripted or based on context. We will also study the impact of trustworthiness between the human and the Cognitive Agent during flight training. We will try to identify the differences in pilot perceptions of trustworthiness, affinity, and preferences in a highly visually and audibly realistic Cognitive Agent. A severe uncanny valley effect may negatively impact training and perceived trustworthiness of the Cognitive Agent. Knowledge extraction and modeling for Cognitive Agents and question answering using NLP will be considered as next steps to convert the raw flight manual or other aviation standard and procedures into the AI solution. We will identify a scalable format and structure for such a knowledge base and ingest historical and generated data set using a large language model (LLM).

ACKNOWLEDGEMENTS
Acknowledgements to Christian Lagarde, Fan E, Sébastien Landry, Jean-Benoit Labbé, Max Hanna, Bill Gunaratnam, Alexandra Ouimet, Dorian Kieken, Yves Paradis, Brian Gariepy, William Duguay, Josée Laganière, Vincent Pilon, Stella-Marissa Hughes, Christopher Courtney and Marc St-Hilaire.

REFERENCES

Brock Lascara, A. L., Matthew DeGarmo, David Maroney, Rick Niles, Lakshmi Vempati. (2019). *Urban Air Mobility Airspace Integration Concepts* (19-00667-9). Retrieved from https://www.mitre.org/sites/default/files/2021-11/pr-19-00667-9-urban-air-mobility-airspace-integration.pdf

CAE. (2021). *Pilot Training for Advanced Air Mobility*. Retrieved from https://www.cae.com/media/documents/Civil_Aviation/cae_evtol_advanced_air_mobility_report_2021.pdf

Charles T. Howell, K. J. S. (2019). *The NASA Urban Air Mobility Testbed Flight Research Aircraft* (NF1676L-32310). Retrieved from Hampton, VA, United States: https://ntrs.nasa.gov/api/citations/20200002700/downloads/20200002700.pdf

De Lellis, E., Bove, E., Ciniglio, U., Corraro, F., Corraro, G., Filippone, E., . . . Palumbo, R. (2019). *Evaluating RPAS Integration with General Aviation Traffic by means of Real-Time Human-in-the-Loop Simulations*. Paper presented at the AIAA Scitech 2019 Forum.

EASA. (2023a). *Artificial Intelligence Roadmap 2.0*. Retrieved from https://www.easa.europa.eu/en/document-library/general-publications/easa-artificial-intelligence-roadmap-20

EASA. (2023b). *EASA Concept Paper: Fisrt usable guidance for level 1 & 2 Machine Learning Applications*. Retrieved from https://www.easa.europa.eu/en/downloads/137631/en

George Price, M. K., Steve Parker, Russell Wolfe. (2020). *Urban Air Mobility Operational Concept (OpsCon) Passenger-Carrying Operations* (NASA/CR—2020–5001587). Retrieved from https://ntrs.nasa.gov/api/citations/20205001587/downloads/UAM%20Passenger-carrying%20OpsCon%20-%20v14%20GP%20accept.pdf

Jeffrey Homola, T. P., Joey Mercer, Nancy Bienert, Conrad Gabriel. (2016). UAS Traffic Management (UTM) Simulation Capabilities and Laboratory Environment. *NASA Ames Research Center*.

Md. Shirajum Munir, S. H. D., Kamrul Hasan, Tariqul Islam, Sachin Shetty. (2023). Artificial Intelligence-Enabled Exploratory Cyber-Physical Safety Analyzer Framework for Civilian Urban Air Mobility. *Applied Sciences, 13*(2). doi:10.3390/app13020755

Sai Krishna Gottipati, S. K., Clodéric Mars, Gregory Szriftgiser, François Chabot. (2021). Cogment: Open Source Framework For Distributed Multi-actor Training, Deployment & Operations. *arXiv:1802.07876*(2106.11345v1).

Shulu Chen, A. E., Marc Brittain and Peng Wei. (2023). Integrated Conflict Management for UAM with Strategic Demand Capacity Balancing and Learning-based Tactical Deconfliction. *arXiv*. Retrieved from https://arxiv.org/abs/2305.10556

Uri Pelli, R. R. (2020). Flying-cab drivers wanted. Retrieved from https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/flying-cab-drivers-wanted

Volodymyr Mnih, A. P. B., Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, Koray Kavukcuoglu. (2016). *Asynchronous Methods for Deep Reinforcement Learning*. Paper presented at the Proceedings of The 33rd International Conference on Machine Learning. https://proceedings.mlr.press/v48/mniha16.html

Xiaoge Zhang, S. M. (2019). Ensemble machine learning models for aviation incident risk prediction. *Decision Support Systems, 116*, 48-63.

Yibing Xie, a. N. P. (2020). *Artificial Intelligence Explanation for Decision Supporting System in Low Level Air Traffic Management*. Paper presented at the Proceedings AIAA Infotech@Aerospace.