

Challenges in Close-Proximity Safe and Seamless Operation of Manned and Unmanned Aircraft in Shared Airspace

Jay Patrikar¹, Joao P. A. Dantas^{1,2}, Sourish Ghosh¹, Parv Kapoor¹, Ian Higgins¹,
Jasmine J. Aloor³, Ingrid Navarro¹, Jimin Sun¹, Ben Stoler¹, Milad Hamidi¹,
Rohan Baijal⁴, Brady Moon¹, Jean Oh¹, Sebastian Scherer¹

Abstract—We propose developing an integrated system to keep autonomous unmanned aircraft safely separated and behave as expected in conjunction with manned traffic. The main goal is to achieve safe manned-unmanned vehicle teaming to improve system performance, have each (robot/human) teammate learn from each other in various aircraft operations, and reduce the manning needs of manned aircraft. The proposed system anticipates and reacts to other aircraft using natural language instructions and can serve as a copilot or operate entirely autonomously. We point out the main technical challenges where improvements on current state-of-the-art are needed to enable Visual Flight Rules to fully autonomous aerial operations, bringing insights to these critical areas. Furthermore, we present an interactive demonstration in a prototypical scenario with one AI pilot and one human pilot sharing the same terminal airspace, interacting with each other using language, and landing safely on the same runway. We also show a demonstration of a vision-only aircraft detection system. [\[Video\]¹](#)

I. INTRODUCTION

Advanced Aerial Mobility (AAM) is an inclusive term that covers urban (UAM), regional (RAM), intraregional (IAM), and suburban air mobility (SAM) solutions [1]. All of these proposed solutions have one thing in common: They all envision a future where autonomous or semi-autonomous aerial vehicles are seamlessly integrated into the current airspace system. AAM solutions open doors to significant socio-economic benefits while at the same time presenting challenges in the seamless integration of these systems with human-operated aircraft and controlling agencies.

Today, manned and unmanned vehicles are separated, limiting the utility and flexibility of operations and reducing efficiency. Human operated aircraft, follow one of the two rules for operation: Visual Flight Rules (VFR) or Instrument Flight Rules (IFR). The choice of flying under VFR or IFR is typically a function of weather conditions. Under

VFR, an aircraft is flown just like driving a car within the FAA established rules of the road. IFR flying, on the other hand, is typically associated with flying aircraft under degraded weather conditions where separation is provided by the ATC. While a pursuit of integrating AAM can start either under IFR or VFR, automated-VFR operations often have better scalability than automated-IFR operations [2]. Another option involves UTM solution [1], which in its current iteration only focuses on small unmanned aerial aircraft operating close to the ground (less than 700 ft) in uncontrolled airspace.

Mastering visual flight rules (VFR) operations for autonomous aircraft has significant operational advantages at unimproved sites, as well as in achievable traffic density compared to instrument flight rules (IFR) or completely separated operations between manned and unmanned systems. For (semi-)autonomous aircraft to operate in tandem with human pilots and ATC controllers under VFR, technological advancements in certain key areas are required, specifically:

- 1) Unmanned aircraft should be able to guarantee self-separation even in a tightly-spaced terminal airspace environment;
- 2) Unmanned vehicles should be able to interpret high-level instructions by ATC to meet the expectations of a normal traffic flow;
- 3) Autonomous aircraft need to understand the expected and unexpected behavior of other aircraft;
- 4) Communications by other pilots and ATC need to be parsed and valid, corresponding responses should be produced; and
- 5) Other aircraft need to be detected and estimated, as well as their future trajectories, need to be predicted.

Many of these challenges have parallels in the self-driving industry, and the technological improvements there can be leveraged to produce domain-specific solutions for AAM. While this is promising, VFR-like AAM integration introduces newer challenges while pushing boundaries on the current state of technology.

II. TECHNICAL CHALLENGES

This section brings the key areas of development where improvements on current state-of-the-art are needed to enable VFR-like autonomous aerial operations. Also, we bring some solutions considering the technical challenges pointed out.

¹Authors are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. {jpatrika, jdantas, sourishg, parvk, ihiggins, ingridn, jimins2, bstoler, mmoghass, bradym, hyaejino, basti}@andrew.cmu.edu

²Joao P. A. Dantas is with the Decision Support Systems Subdivision, Institute for Advanced Studies, Sao Jose dos Campos, SP, Brazil. dantasjpad@fab.mil.br

³Jasmine J. Aloor is with the Indian Institute of Technology, Kharagpur, WB, India. jasminejerry@iitkgp.ac.in

⁴Rohan Baijal is with the Indian Institute of Technology, Kanpur, UP, India. rbaijal@iitk.ac.in

¹Video: https://youtu.be/iU_MyMwuE8E

A. Aircraft Detection and Tracking

See-and-Avoid is one of the key tenets of VFR operation. The ability to spot other aircraft or aerial hazards like birds, balloons, gliders, etc., and execute maneuvers so as to mitigate a collision hazard is critical to successfully deploying AAM solutions in the NAS. Detecting aircraft at long ranges comes with a lot of challenges. Some of these are outlined as follows:

- **Low SnR:** Aircraft or other flying objects at long ranges appear very small and typically have very low signal-to-noise (SnR) ratios (see Fig. 1)
- **Poor performance on small objects:** The state-of-the-art in object detection keeps changing ever so frequently, and the standard benchmarks such as COCO [3] are always being beaten from time to time as more complex models are created. However, the performance of these detectors on small objects ($< 32^2$ px²) is pretty poor as they have low average precision and recall. In the context of our work, the performance is even worse since most of our data consists of small objects.
- **Computational constraints:** The SOTA detectors are typically very complex in terms of the number of learnable parameters and they usually work on low to moderate image resolutions. Since our context is primarily small objects, we need to operate on high-res 2K images. Such images cannot be processed without downsampling by modern detectors.



Fig. 1. Visualization of the detected aircraft bounding boxes from the data onboard a general aviation aircraft. At long ranges, the SnR is really poor, and our algorithms are trained to detect aircraft even in low SnR situations.

B. Intent Prediction

Reasoning about the potential set of future trajectories that other aircraft can take is critical to ensure that the best actions are taken that also minimize risk. Typical prediction methods use short horizons and fail to capture the long-term intent of other aircraft. The majority of trajectory prediction work has been explored in the pedestrian and autonomous vehicle domains. Within AAM, goal points such as airports, pilots, ATC communications, and rules of way (such as FAR §91.113) are often known. The use of this explicit source of information as well as implicit sources like weather can help decipher the intent of other aircraft and increase the length of reliable predictions.

The lack of datasets and baseline methods makes it difficult to conduct meaningful research. Towards this end,

we have been collecting traffic data from transponders at select controlled and uncontrolled airports to understand and train models for intent prediction. To capture the influence of weather on pilot decisions, the weather data was also collected by parsing Meteorological Aerodrome Reports (METARs) to gather the relevant sections from the full weather report. We recently published the first chunk of the trajectory data TrajAir along with a baseline method christened TrajAirNet to predict aircraft trajectories trained on the collected data (see Fig. 2) [4].

C. Safe and Socially-Compliant Navigation

As the FAR rules only specify a rough guideline, autonomous vehicles must be equipped with the capability to make flexible decisions to comply with the traffic norm of arbitrary situations. The idea of social navigation is to learn to follow the observed traffic patterns in a socially-compliant manner. Generating actions that are not only safe but also socially compliant while following rules is thus critical in generating behavior that is acceptable to human pilots co-habiting the same airspace. We use variants of Monte Carlo Tree Search algorithms that combine learned behavior with logic specifications like Signal Temporal Logic to enable safe and socially compliant navigation (see Fig. 3).

D. Guaranteeing safety

The safety system ensures that the future trajectory of the aircraft is always safe and uses projected distance maximization to ensure collision-free flight. The system performs a minimally invasive modification to the plans made by the planning/inference engine to ensure that these plans will not violate safety invariants. These safety invariants are primarily focused on avoiding collisions with other agents and risky maneuvers. The system assumes linear velocity vectors for other agents while interpolating for the future. The system is also designed with runtime performance objectives in mind and ensures safe behavior on the fly.

E. Automated speech recognition and production

It is critical to establish clear communication between a human operator/pilot and an AI system in our target problem domain. Specific to the aviation domain is the challenge of understanding and decoding aviation-specific terminology that is different from everyday speech constructions. Other challenges such as radio background noise, incomplete instructions, and radio phraseology also need to be addressed.

We propose developing a bi-directional communication mechanism for human-AI collaboration, focusing on clarity instead of naturalness to accomplish an acceptable performance level to produce a language covering the controlled vocabulary used in airspace operations. We employed learning approaches for the AI system to understand complex concepts, e.g., learning from demonstrations, and, to language understanding, developed a language generation system, which can summarize visually perceived information [5]. This language generation capability can also clarify potential ambiguity when receiving commands from a human

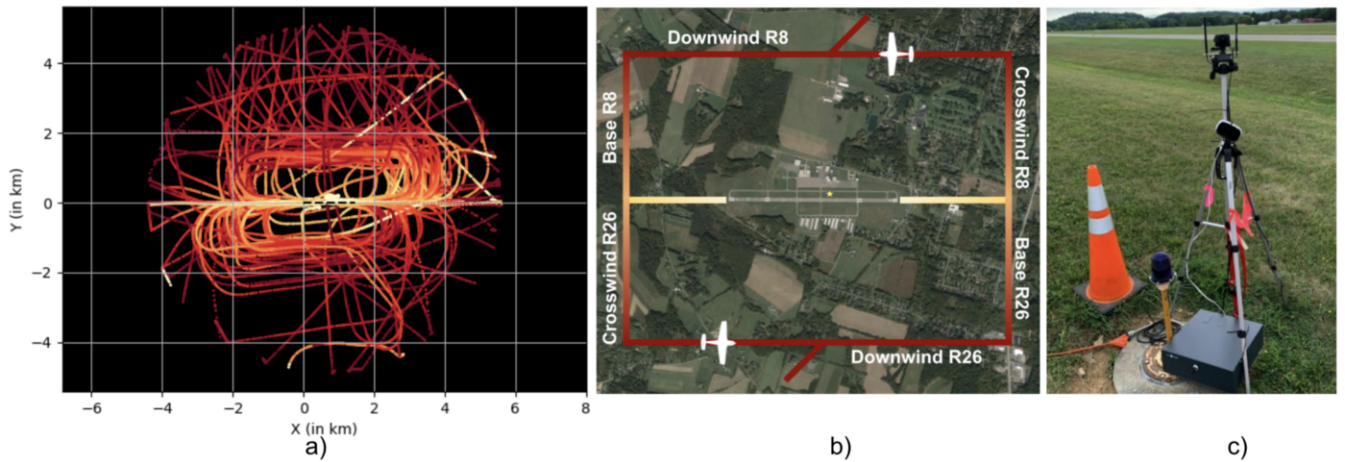


Fig. 2. The figure [4] shows the dataset and its collection setup at the Pittsburgh-Butler Regional Airport (KBTP) — a non-towered GA airport that serves as a primary location for the dataset. Lighter colour indicates lower altitude. a) Shown is a snippet of the processed dataset with aircraft trajectories showing clear lobes for traffic patterns for both runways. b) The left traffic pattern and nomenclature for the runways at the airport. c) Picture of the data-collection setup.

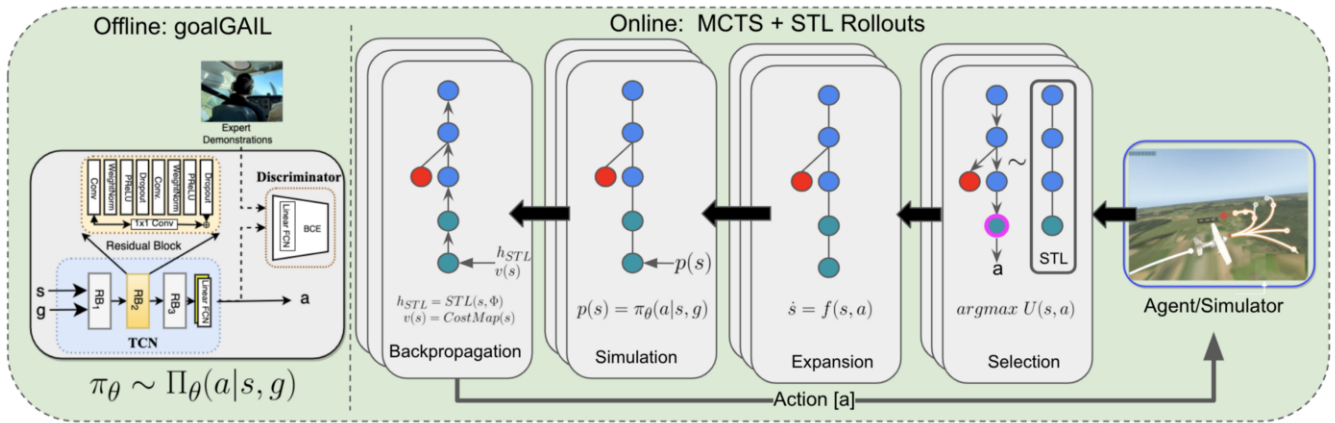


Fig. 3. Figure shows the proposed planning setup for a single agent. The offline policy is trained using a Generative Adversarial Imitation Learning algorithm that learns system behavior from the recorded trajectories from KBTP airport. Online, this behavior is executed using a Monte Carlo Tree Search algorithm that uses multiple roll-outs to model possible future states of the human agent. The roll-outs are constrained using Signal Temporal Logic specifications that ensure that the actions follow the rules established by the FAA.

operator/pilot. To build a language representation model customized for aviation, we will leverage a large corpus of aviation relevant documents such as manuals to train a model, and, to use such unlabeled data, we employed a transformer-based BERT [6] variant to pre-train from the unlabeled text.

F. High-fidelity Simulators

For safety-critical domains, accurate high-fidelity simulators are required to test algorithms before real-world testing. Given the lack of simulators in the public domain, we designed our simulator christened XPlaneROS. With X-Plane 11 as the high-fidelity flight simulator and ROSplane [7] as the autopilot, we obtain realistic aircraft models and visuals of similar aircraft (see Fig. 4). The X-Plane Connect Toolbox interfaces between X-Plane and Robot Operating System (ROS) topics. Based on the high-level commands input by the planner, ROSplane generates the actuator commands for

ailerons, rudder, elevator, and throttle, which are in turn sent to X-Plane through XPlaneConnect. ROSplane uses a cascaded control structure and has the ability to follow waypoints with Dubin’s Paths. XPlaneROS provides additional capabilities to follow a select set of motion primitives. There have also been some extensions to ROSplane, like employing a proper takeoff, additional control loops for vertical velocity rates, and a rudimentary autonomous landing sequence.

III. INTERACTIVE DEMONSTRATION SUMMARY

For the interactive demonstration, we propose a scenario with one AI pilot and one human-pilot landing at the same runway in a non-towered airfield with runways 08 and 26. The goal of the AI system is to execute a low approach over the Runway 26, use a natural language process algorithm to understand the human intent and coordinate with the human pilot to achieve the joint objective (see Fig. 5). The interactive demonstration has the following stages:

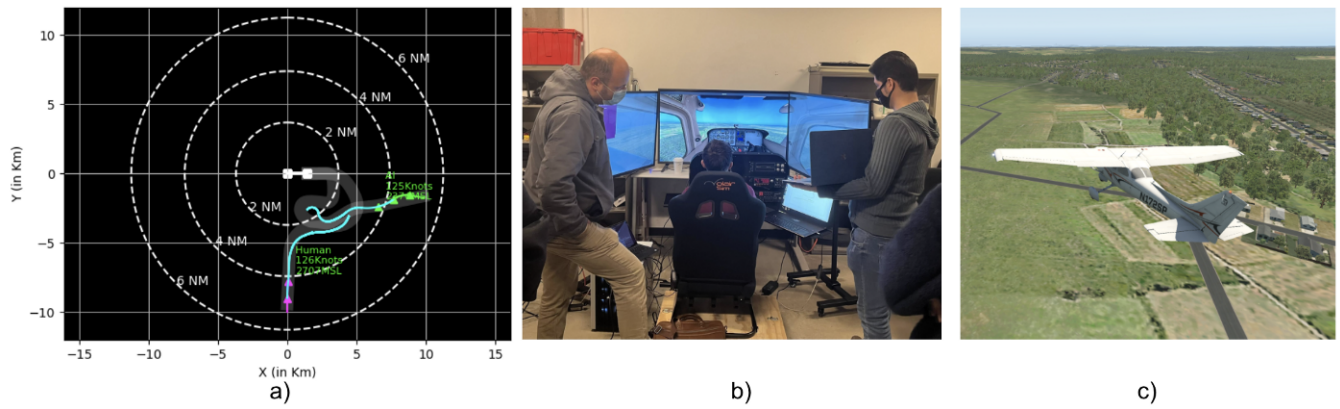


Fig. 4. The figure shows the high-fidelity simulator setup that enables Human-AI interaction. Figure a) shows a top-down view with one human agent (magenta) interacting and one AI agent (lime) while trying to land on the same runway. The most likely branch of the MCTS forward propagation tree for both the agents is shown in cyan. White lines show the reference trajectories. Figure b) shows the physical simulator setup with an immersive environment for the human pilot. Figure c) shows a screengrab of the visual rendering of the simulator using the X-Plane 11 flight simulator backend.

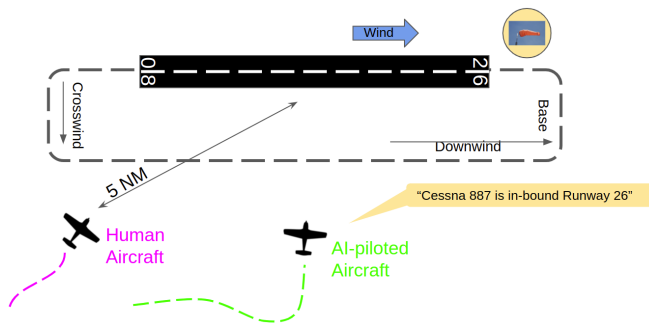


Fig. 5. Scenario depiction of the interactive demonstration.

- 1) The human pilot is erroneously heading to land on runway 08 and informs their intention on a broadcast message.
- 2) The AI pilot generates a safe plan by predicting the human pilot's intention from the language by listening to the human pilot's broadcast message.
- 3) The AI pilot correctly calculates that the preferred runway is 26 and generates speech to broadcast its goal of landing on runway 26 to the human pilot.
- 4) The human pilot decided to change to runway 26 and inform this new intention on a new broadcast message.
- 5) The AI pilot generates the new plan changing to runway 26. Once more, the AI pilot generates speech to broadcast its new goal to the human pilot.
- 6) Both the AI and the human pilots, knowing the position of each other, coordinate the landing on runway 26.

A separate demonstration of vision-only long-range aircraft detection is also showcased using real-world data.

IV. CONCLUSIONS

We believe that the current understanding of integrating unmanned aircraft within the National Airspace System (NAS) needs a more human-pilot-centered approach. Pilots and aircraft already operating in the national airspace are essential stakeholders in the larger discourse.

Technological developments enabling close-proximity safe and seamless operation of manned and unmanned aircraft in a shared airspace need an understanding of the rules and conventions already in place. The NAS is a dynamic environment with in-built flexibility and protocols to handle traffic volumes and emergencies. Our core understanding is that rather than having the current NAS adapt to changing autonomy needs, we need to move towards identifying technological requirements that enable unmanned systems to operate along with humans collaboratively. We hope to build true AI pilots indistinguishable from human pilots to allow seamless integration within the current NAS.

ACKNOWLEDGMENT

This work is supported by the Army Futures Command Artificial Intelligence Integration Center (AI2C).

REFERENCES

- [1] A. S. Aweiss, B. D. Owens, J. Rios, J. R. Homola, and C. P. Mohlenbrink, *Unmanned Aircraft Systems (UAS) Traffic Management (UTM) National Campaign II*. Aerospace Research Center, 2018. [Online]. Available: <https://arc.aiaa.org/doi/abs/10.2514/6.2018-1727>
- [2] E. R. Mueller, P. H. Kopardekar, and K. H. Goodrich, "Enabling airspace integration for high-density on-demand mobility operations," in *17th AIAA Aviation Technology, Integration, and Operations Conference*, 2017, p. 3086.
- [3] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [4] J. Patrikar, B. Moon, J. Oh, and S. Scherer, "Predicting like a pilot: Dataset and method to predict socially-aware aircraft trajectories in non-towered terminal airspace," 2021. [Online]. Available: <https://arxiv.org/abs/2109.15158>
- [5] J. Tian and J. Oh, "Image captioning with compositional neural module networks," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 7 2019, pp. 3576–3584. [Online]. Available: <https://doi.org/10.24963/ijcai.2019/496>
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [7] G. Ellingson and T. McLain, "Rosplane: Fixed-wing autopilot for education and research," in *Unmanned Aircraft Systems (ICUAS), 2017 International Conference on*. IEEE, 2017.