

AI Pilot: Close-Proximity Safe and Seamless Operation of Manned and Unmanned Aircraft in Shared Airspace

1. Objectives: Develop an AI system to keep autonomous unmanned aircraft, in conjunction with manned traffic safely separated and behave as expected when entering and leaving the traffic/break/formation pattern.

2. Key Areas:

2.1. Intent Prediction

Understanding and predicting the intended motion of humans in an environment is an important skill that robots across various domains, such as aerial robotics, must be equipped with in order to enable safe interactions.

Motion and intent are influenced by several factors, making the task challenging. One factor is social behavior, which describes how agents interact, and it heavily depends on the context. Within AAM, goal points such as airports, pilots, ATC communications, and rules of air are often well-known. Using this definitive source of information and other implicit sources like weather can help decipher the intent of other aircraft and increase the length of reliable predictions. Respecting motion constraints is another relevant aspect in this setting, as it is related to restrictions that may arise from the agent's own physical constraints, or the constraints imposed by the environment, such as the topology of a scene, or the rules associated with it, *e.g.*, pilots should respect flying guidelines. Moreover, agent motion is driven by each agent's goals which vary depending on the situation and the type of agent: goals can be flexible or fixed, short or long-term, implicit or explicit. This varied nature makes them difficult to model. The majority of existing works on trajectory prediction are mainly focused on pedestrian behaviors where the context is relatively loosely defined, *i.e.*, except for the social norm that pedestrians try to maintain a comfortable distance from others, there are few rules that guide pedestrian behavior.

2.2. Planning

As the FAR rules only specify a rough guideline, autonomous vehicles must be equipped with the capability to make flexible decisions to comply with traffic norms and generalize to arbitrary situations. The idea of rule-based navigation is to learn to follow the observed traffic patterns. Generating actions that are not only safe but also follow rules is thus critical in generating behavior that is acceptable to human pilots co-habiting the same airspace.

Deep reinforcement learning methods have been successful in learning autonomous navigation. In sequential decision-making, policies are often represented by a Markov Decision Process

(MDP) with a well-defined reward function. The complexity of manually defining a reward function hampers the widespread applicability of reinforcement learning and optimal control algorithms. Learning a policy from expert demonstrations (LfD or imitation learning) has led to a variety of applications with state-of-the-art performance. Behavior cloning (BC) enables the implementation of supervised learning like methods but requires a large quantity of demonstration data and interactive learning to reduce compounding errors due to out-of-distribution data or covariate shift. These approaches, however, have the limitations of being sensitive to sub-optimal demonstrations and the inflexibility to generalize to other applications. Generative adversarial imitation learning (GAIL) and related works have approached imitation learning as a system consisting of a learner's policy and a discriminator network that attempts to distinguish between learner transitions from expert transitions. GAIL attempts to learn policies directly from simulation and has the limitations of sub-optimal demonstrations, and convergence becomes slow due to stability issues with low data.

Temporal logics such as Signal Temporal Logic (STL) provide a mathematically robust representation to encode such spatio-temporal constraints. They can be used to logically specify desired behavior translated from requirements expressed in natural language. STL is used to specify properties over real-valued dense time signals often generated by continuous dynamical systems. Quantitative semantics associated with STL provides a real value called robustness which quantifies the degree of satisfaction or violation. This property enables the designer to encode domain-specific constraints and quantitatively measure their satisfaction with motion planning and control. In order to infuse the STL specifications to improve the base LfD, we propose using a variant of Monte Carlo Tree Search (MCTS) that uses an offline pre-trained network to generate simulated roll-outs. MCTS is a powerful heuristic search algorithm often deployed for long-horizon decision making tasks.. MCTS, when used with a UCT heuristic, has properties like anytime convergence to the best action. This makes it well suited to be used as a long-horizon goal directed planner within large state spaces with a time budget.


2.3. Automated speech recognition and production

Establishing clear communication between a human operator/pilot and an AI system in our target problem domain is critical. Understanding and decoding aviation-specific terminology, which differs from everyday speech constructions, is a big challenge. Other challenges such as radio background noise, incomplete instructions, and radio phraseology must also be addressed.

To this end, we propose 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 employ learning approaches for the AI system to understand complex concepts, e.g., learning from demonstrations, and, to language understanding, develop a language generation system, which can summarize visually perceived information. This language generation capability can also clarify potential ambiguity when receiving commands from a human operator/pilot.

2.4. High-fidelity Simulator and Integration

A simulator that demonstrates a specific performance level and satisfies some standards may be applied to training, testing, and checking instead of doing the same during flight. The aviation-training community believes that a high level of fidelity is mandated to deliver the highest level of transfer of learning to the actual equipment. The fidelity of a flight simulator is essential to the pilots' performance in a real world. We are constantly working to deliver high-fidelity simulation-based assessment devices to enhance the expert pilot performance in real-world situations. Therefore, simulations have been implemented to help aviators refine their skills to become exceptional performers in the aviation industry.



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.

3. Technical Results

3.1. Intent Prediction




Figure shows the proposed intent prediction model.

For this project, we specifically focus on the settings where there exist navigation guidelines relatively strictly defined yet not clearly marked in their physical environments, e.g., in aerial navigation, although there are strict rules analogous to ground navigation, the air space does not display visible lanes. To address this challenge, we propose an approach, known here as Social-PatteRNN [1], where we aim to learn such contexts in the form of motion patterns and social influences extracted from the patterns. **This work builds on background IP “Social-PatteRNN: An algorithm for socially-aware trajectory prediction from motion patterns” (under submission).** We hypothesize that, within these domains, in the short term, agents tend to exhibit some motion patterns that reveal both their general directions of motion and their intermediate goals. We also believe that such patterns further explain general rules of motion, which may be explicit, e.g., sports rules or flying guidelines, or implicit, e.g., social etiquette, as well as, capture admissible motions. From these intuitions, we propose a data-driven approach to learn these patterns of motion and use them as a conditioning signal for predicting multimodal trajectories. Then, we use the learned patterns to extract sub-goal information which we aggregate to our model as the social context.

We also assessed the generalizability of the proposed approach across different domains: humans in crowds, humans in sports, and manned aircraft. Experimental results show that the proposed approach has consistent performance across these domains.

Table shows the results for various intent prediction algorithms compared to ours. Results are in terms of Average Displacement Error and Final Displacement Error ADE / FDE.

Trajair (km)		
1	A-VRNN [5]	0.64 / 1.31
2	DAG-Net [5]	0.78 / 1.53
3	TrajAirNet [3]	0.77 / 1.50
4	S-PEC [10]	0.96 / 2.05
5	Social-PatteRNN (Ours)	0.56 / 1.20
6	Social-PatteRNN-ATT (Ours)	0.55 / 1.19




Figure shows various models' intent prediction results compared against our model, *Social-PatteRNN*. Dashed lines represent the ground truth trajectory, while solid lines represent the predicted intent. The proposed approach exhibits less jerkiness, producing more stable trajectories and closer to the ground truth, compared to other models.

3.3 Planning




Figure shows the proposed planning algorithm

In this work, we present a novel decision-making method that uses MCTS to build a tree structure that guides the offline pre-trained LfD policies online with STL specifications. We achieve this by biasing the MCTS heuristic sampling towards branches with higher STL satisfactions. The primary insight is that while the LfD policy decides on the low level executions in line with the expert demonstrations, STL encourages the satisfaction of high-level objectives. This hierarchical approach provides a method to encode rules while allowing the agent to choose how to satisfy the constraints in a learning-enabled framework. The STL specifications thus provide a guide rail against the low-bandwidth noise in the expert demonstrations.

Given a continuous-space dynamic system of the form $\dot{s} = f(s, a)$, we define a discrete-time Markov Decision Process without rewards, (MDP \mathbb{R}). Let $M = (S, A, T, \rho_0, G)$ where S is the set of states $s \in S$, $a \in A$ is the discrete set of actions or motion primitives that follow $f(\cdot)$, $T : S \times A \Rightarrow S$ is the transition function, $\rho_0 \in S$ is an initial state distribution, and G is the goal distribution. The task is to produce a policy $\pi(\theta)$ from a start location $s_0 \in \rho_0$ to goal location $g \in G$ that leads to a trajectory $\tau = (s_0, a_0, s_1, a_1, \dots, g)$. We also assume access to expert trajectories $D = \{(s_0, a_0, s_1, a_1, \dots, g)\}$ and high-level STL specification Φ that encodes any rules we expect the system to follow. An STL formula Φ can be built recursively from predicates using the following grammar

$$\Phi := \top \mid \mu c \mid \neg \Phi \mid \Phi \wedge \Psi \mid \diamond [a, b] \Phi \mid \square [a, b] \Phi \mid \Phi_1 U [a, b] \Phi_2$$

where Φ_1, Φ_2 are STL formulas, \top is the Boolean True, μc is a predicate of the form $\mu(s) > c$, \neg and \wedge the Boolean negation and AND operators, respectively, and $0 \leq a \leq b < \infty$ denote time intervals. The temporal operators \diamond , \square and U are called “eventually”, “always”, and “until” respectively. The quantitative semantics of a formula with respect to a signal $\square xt$ can be used to compute robustness values for the specifications used in our application.

Framework

Algorithm 1: Plan (θ, Φ)

```

1  $s \leftarrow \text{Sample}(\rho_0)$ 
2  $g \leftarrow \text{Sample}(G)$ 
3 while  $s \neq g$  and not  $\text{maxSteps}$  do
4   while  $\text{timeElapsed} \leq \text{planHorizon}$  do
5     |  $N(\cdot) \leftarrow \text{MCTS}(s_{0:t}, g, \theta, \Phi, 0)$ 
6   end
7    $a \leftarrow \text{choice}_{a'}(N(s, a'))$ 
8    $s \leftarrow T(s, a)$ 
9 end

```

We first train a LfD policy by formulating the problem as finding a distribution of future actions conditioned on the past trajectories and the goal where tobs is the observation time horizon. The MCTS (see Algorithm 2) uses the policy π^θ to generate simulations. Each simulation starts from the root state and iteratively selects moves that maximize the STL modified UCT heuristic. For each state transition, we maintain a directed edge in the tree with an action value $Q(s, a)$, prior probability $P(s, a)$, STL heuristic $H(s, a)$ and a visit count $N(s, a)$. The total heuristic value is calculated as a weighted sum by controlling the degree of exploration and STL heuristic's weight. Starting with the initial state $s(0)$, at each time step, we calculate the action to take, which maximizes $U(s, a)$ (Line 12). If the next state already exists in the tree, we continue our simulation, else a new node is created in our tree, and we initialize its $P(s, \cdot) = \square p\theta(s)$ from our policy π^θ (Line 8). The expected reward $v = v\theta(s)$ can be provided by the user as a learned value function or as a cost-map (Line 9). The heuristic h_{STL} is calculated using the STL specification (Line 14). We initialize $Q(s, a)$, $H(s, a)$ and $N(s, a)$ to 0 for all a . We then propagate the cost v and the STL heuristic h_{STL} back up the MCTS tree, updating all the $Q(s, a)$ and $H(s, a)$ values seen during the simulation, and start again from the root.

After running forward simulations of the MCTS, the $N(s, a)$ values provide a good approximation for the optimal stochastic process from each state. Hence, the action we take is randomly sampled from a distribution of actions with probability proportional to $N(s, a)$. We expand the search space until we exhaust the planning budget time planHorizon . After each action is selected, the MCTS tree is reinitialized from the actual trajectory followed by the agent. The planner is terminated when the goal is reached, or the maximum number of steps is reached, whichever comes first.

Algorithm 2: MCTS ($s_{0:t}, g, \theta, \Phi, h_{STL}$)

```

1 if  $s \in G$  then
2   | return  $s == g, h_{STL}$ 
3 end
4 if  $s \notin Tree$  then
5   |  $Tree \leftarrow Tree \cup s$ 
6   |  $Q(s, a) \leftarrow 0$ 
7   |  $N(s, a) \leftarrow 0$ 
8   |  $P(s, a) \leftarrow \hat{\pi}_\theta(s)$ 
9   |  $v(s) \leftarrow CostMap(s)$ 
10  | return  $v(s), h_{STL}$ 
11 else
12  |  $a \leftarrow argmax_{a'} \left[ Q(s, a') + \frac{c_1 P(s, a') \sqrt{N(s)}}{1 + N(s, a')} + \right.$ 
13  |  $c_2 H(s, a') \left] \right.$ 
14  |  $s' \leftarrow T(s, a)$ 
15  |  $h_{STL} \leftarrow STL(s' + s_{0:t}, \Phi)$ 
16  |  $v, h_{STL} \leftarrow MCTS(s' + s_{0:t}, g, \theta, \Phi, h_{STL})$ 
17  |  $N(s, a) \leftarrow N(s, a) + 1$ 
18  |  $Q(s, a) \leftarrow \frac{N(s, a) Q(s, a) + v}{1 + N(s, a)}$ 
19  |  $H(s, a) \leftarrow \frac{N(s, a) * H(s, a) + h_{STL}}{1 + N(s, a)}$ 
20 end

```

Implementation Details

Consider a fixed-wing aircraft at time t , let $s_t = (x_t, y_t, z_t, \chi_t) \in \mathbb{R}^3 \times SO(2)$ denote the position and where v is the aircraft's inertial speed, v_{2D} is the speed in the 2-D plane, ϕ is the roll-angle, and v_h is vertical speed. Finally, g is the acceleration due to gravity. We assume a zero wind condition.

The action space A is a fixed library of 30 motion primitives that discretizes each of the control inputs. We use the inertial velocities (v) 70 and 90 knots, the vertical velocities (v_h) +500 ft/min and -500 ft/min and the bank angle (ϕ) discretization such that χ changes by 45° and 90° heading over the chosen time-horizon of 20 sec. The goal distribution G is a one-hot vector representation of the final goal of a particular agent as the eight cardinal directions along with two runway ends, as shown in Fig. 3. The set G is represented as $G = \{N, NE, E, SE, S, SW, W, R08, R26\}$ with each element representing the final region the aircraft is desired to reach as shown in Fig. 3. For simplicity we also set the start states equal to G .

The implementation details are split between online and offline components.

- 1) Dataset: We use the TrajAir dataset (5.2.1), which is collected at the Pittsburgh-Butler Regional Airport (ICAO:KBTP). The dataset contains 111 days of transponder data that can be

used to extract expert demonstrations from pilots as they navigate the un-towered airspace. The dataset trajectories are smoothed using a B spline (basis-spline) approximate representation of order 2.

2) Offline LfD Policy Details: The LfD policy takes as input the past trajectories of the agent to predict its possible action distribution. While the method can use any LfD policy, we use a goal-conditioned generative adversarial imitation learning (GoalGAIL) method. The GoalGAIL is modified to use Temporal Convolutional Layers (TCNs) to process the sequential trajectory data. TCN layers encode a trajectory's spatio-temporal information into a latent vector without losing the underlying data's temporal (causal) relations. We use TCNs as an alternative to using LSTMs for encoding the trajectories.

We break the trajectories in a scene into sequences of length $t_{obs} + t_{pred}$ where $t_{obs} = 11\text{sec}$ and $t_{pred} = 20\text{sec}$. In a given scene, the raw trajectory in absolute coordinates of the agent is encoded using the TCN layers as h_{obs} . The agent's goal $g \sqsubseteq G$ is encoded through an MLP layer, φ_1 , and is concatenated with the encoded trajectory vector. We measure how close the predicted trajectory is to the expert trajectory using a mean squared error (MSE) loss. A discriminator $D\psi$ is trained to distinguish expert transitions from policy transitions. The combination of these two loss functions is used to train the model.

In order to convert s^* to a^* , we match the generated trajectories from the control inputs in the library A using a weighted L2 Euclidean error distance over (x, y, z) points on the trajectory. For training, we use the AdamW optimizer with a learning rate of $3e - 3$.

3) Signal Temporal Logic Specifications: We evaluate the performance of our agent based on reaching the goal while adhering to the airport traffic pattern. The goal objective, as well as traffic pattern compliance, are both encoded using STL specifications. We use the three stages for landing pattern. Φ_1 , Φ_2 , and Φ_3 represent the STL formulas encoding occupancy of regions corresponding to the downward, base, and final stages, respectively. The landing STL specification becomes:

$$\Phi_L = \diamond(\Phi_1 \wedge \diamond(\Phi_2 \wedge \diamond\Box\Phi_3)))$$

$\diamond(\Phi)$ can be interpreted as "Eventually" being in a region represented by Φ . The nested \diamond operators encode a sequential visit of regions represented by Φ_1 , Φ_2 , and Φ_3 . Similarly, the takeoff STL specification is defined based on the goal regions reached by the aircraft. By defining reaching a goal region $g \sqsubseteq G$ by an STL formula Φ_4 , we get the takeoff specification:

$$\Phi_T = \diamond(\Phi_4)$$

The robustness values of the STL specifications are evaluated on the state trajectory traces generated by the search tree. The first element of the traces is the tree root node, and the last element of the trace is the node whose value is being computed.

4) Online Monte Carlo Tree Search : The MCTS is implemented as a recursive function where each iteration ends with a new leaf that corresponds to an action in the trajectory library. The implementation uses a normalized costmap $v(s)$ that is built by counting the frequencies of the aircraft in the TrajAir dataset at particular states s .

Evaluations

Evaluation of the proposed approach is performed using a custom simulator that follows the dynamics defined in Eq. 4. The network implementations are in PyTorch. We use the recently released rtamt package, an online monitoring library for calculating STL robustness values. To showcase real-time online evaluations, simulations are performed on an Intel NUC computer with Intel® Core™ i7- 8559U CPU @ 2.70GHz × 8. The complete implementation details and parameter details are in the open-sourced code base3and the associated Readme.

A. Qualitative results




Figure shows an example scenario where the aircraft starts from the South-West and is tasked with landing at R26 while following the standard FAA traffic pattern. The rules of the traffic pattern are encoded as STL specifications. The white marked line shows the aircraft's position at each step. At every step, the MCTS replans, and the resulting tree is shown in magenta. The STL sub-specifications are shown as rectangles. As can be seen, the aircraft manages to reach the runway while satisfying the specifications. The size of the search tree is a function of available planHorizon.

B. Comparative results

In order to perform quantitative studies, we compare the performance of the proposed algorithm with vanilla LfD policies. We uniformly sample 100 start-goal pairs randomly from p_0 , G . G is truncated to N, S, E, W, R where R represents both R08 and R26 to condense the results. We then provide these to Algorithm 1, which plans for the agent.

Based on the selected start and goal pair, a relevant STL specification is chosen. Each episode ends when the agent reaches the goal or if the maximum steps are exceeded. In addition to the goalGAIL policy, we also train a pure Behavior Cloning (BC) policy. The BC policy uses a TCN to encode the history and outputs an action without a goal vector or a discriminator.

Comparisons were carried out with both these policies integrated into the MCTS framework with ablation on the STL heuristic to show the impact of the STL specification.

We define our evaluation metrics as follows:

- Success Rate: Fraction of episodes that were successful in reaching their goal locations.
- STL Score: Average of the normalized fraction of the STL robustness value satisfied over all the episodes. A higher value indicates better satisfaction.

Table I shows the quantitative results. Our proposed algorithm performs significantly better than the baselines for all start-goal pairs. We get an almost perfect success rate in the aircraft takeoff scenarios. The aircraft landing cases are more challenging due to following the specific landing patterns when incoming from different sides of the runway, which is reflected in the success rates shown. Additionally, we observe the baseline BC performs similarly to GAIL on the success metric but not on the STL robustness. This shows that while BC policies are comparable in reaching the goals, the transient performance of GoalGAIL is better.

Incorporating STL improves the robustness values for both LfD policies.

Algorithm	Takeoff Specification Φ_T				Landing Specification Φ_L				Total
	N	S	E	W	N	S	E	W	
BC + MCTS	0.2 / 0.2	0.0 / 0.1	0.3 / 0.1	0.3 / 0.1	0.0 / 0.0	0.1 / 0.4	0.1 / 0.4	0.3 / 0.4	0.1 / 0.2
GoalGAIL + MCTS	0.0 / 0.3	0.1 / 0.4	0.0 / 0.3	0.0 / 0.4	0.3 / 0.6	0.1 / 0.2	0.3 / 0.7	0.1 / 0.5	0.1 / 0.4
BC + MCTS + STL	1.0 / 0.9	1.0 / 0.9	0.6 / 0.3	0.7 / 0.4	0.2 / 0.5	0.2 / 0.5	0.2 / 0.5	0.2 / 0.5	0.5 / 0.6
GoalGAIL + MCTS+STL	1.0 / 0.9	1.0 / 0.9	0.8 / 0.7	1.0 / 0.9	0.7 / 0.9	0.8 / 0.9	0.3 / 0.8	0.5 / 0.9	0.7 / 0.8

TABLE I: Table shows the quantitative results with randomly sampled start and goal states for two LfD policies with ablation studies with the STL heuristic. Results show the Success Rate \uparrow / STL Score \uparrow for two vanilla LfD policies and their ablations with the STL heuristic. Results show both Landing $X \Rightarrow R$ and Takeoff $R \Rightarrow X$ scenarios for each of the cardinal directions X .

3.4 Speech Recognition

We introduce two off-the-shelf modules to realize the automatic speech recognition and speech generation functionality. For speech-to-text, we use a Speech to Text Transformer model trained on the Librispeech corpus. This model achieves a Word Error Rate of 41.3% and a Character Error Rate of 12.9% on our in-house dataset. To further improve the model’s prediction accuracy, we use string matching to replace the predicted words or phrases to those more probable in the air traffic control scenario (e.g., “one way two six” \rightarrow “runway two six”).

Additionally, to investigate the capability of the speech-to-text model even further, we evaluate the model without further fine-tuning on a subset of the LDC Air Traffic Control Corpus, which is a dataset collected from actual pilot utterance recordings. Speech collected in the Air Traffic Control domain typically contains excessive background noise coming from aircraft, and speakers tend to talk much quicker than the Librispeech and Commonvoice corpora general speech-to-text models are trained on. As expected, in this dataset, the model achieves a Word Error Rate of 91% and a Character Error Rate of 60.6%. The relatively high error rate can be explained by the distinct properties of speech recorded in Air Traffic Control domains, and expect that in-domain fine-tuning can boost the recognition accuracy in relevant datasets. For text-to-speech, we employ the recently proposed FastSpeech 2 model and apply slot-filling to a pre-defined set of template utterances.

4. Significant Hardware or Software Developed:

4.1 High Fidelity Simulator




Figure. XPlaneROS Architecture

XPlaneROS (5.2.2) integrates a high-fidelity simulator with a state-of-the-art autopilot. The complete system enables the use of high-level or lower-level commands to control a general aviation aircraft in realistic world scenarios anywhere in the world. We chose X-Plane 11 as our simulator because of its open API and realistic aircraft models and visuals. For the lower-level control, we've integrated ROSplane as the autopilot. ROSplane is a control stack for fixed-wing aircraft developed by the BYU MAGICC Lab.

XPlaneROS interfaces with XPlane 11 using NASA's XPlaneConnect. With XPlaneROS, the information from XPlane is published over ROS topics. The ROSplane integration then uses this information to generate actuator commands for ailerons, rudder, elevator, and throttle based on higher-level input to the system. These actuator commands are then sent to XPlane 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.

5. Publications this Period:

5.1 Papers:

[1] Ingrid Navarro & Jean Oh. *Social-PatteRNN: Socially-Aware Trajectory Prediction Guided by Motion Patterns*. In proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022 (to appear) [\[ArXiv\]](#) .

[2] Aloor, J. J., Patrikar, J., Kapoor, P., Oh, J., & Scherer, S. (2022). Follow The Rules: Online Signal Temporal Logic Tree Search for Guided Imitation Learning in Stochastic Domains. arXiv preprint arXiv:2209.13737.[\[ArXiv\]](#)

[3] Ghosh, S., Patrikar, J., Moon, B., & Hamidi, M. M. (2022). AirTrack: Onboard Deep Learning Framework for Long-Range Aircraft Detection and Tracking. arXiv preprint arXiv:2209.12849.[\[ArXiv\]](#)

[4] Patrikar, J., Moon, B., Oh, J., & Scherer, S. (2022, May). Predicting like a pilot: Dataset and method to predict socially-aware aircraft trajectories in non-towered terminal airspace. In 2022 International Conference on Robotics and Automation (ICRA) (pp. 2525-2531). IEEE.

[5] Patrikar, J., Dantas, J. P., Ghosh, S., Kapoor, P., Higgins, I., Aloor, J. J., ... & Scherer, S. Challenges in Close-Proximity Safe and Seamless Operation of Manned and Unmanned Aircraft in Shared Airspace. Workshop, In 2022 International Conference on Robotics and Automation (ICRA). IEEE.

5.2 Dataset and Software:

[1] Patrikar, Jay; Moon, Brady; Ghosh, Sourish; Oh, Jean; Scherer, Sebastian (2021): TrajAir: A General Aviation Trajectory Dataset. Carnegie Mellon University. Dataset.

<https://doi.org/10.1184/R1/14866251.v1>

[2] Baijal, Rohan; Patrikar, Jay; Moon, Brady; Scherer, Sebastian; Oh, Jean (2021): XPlaneROS : ROS Wrapper for Autonomous Fixed Wing Applications. Carnegie Mellon University. Software.

<https://doi.org/10.1184/R1/16589924.v1>

5.3 Videos:

- [1] STL-MCTS <https://www.youtube.com/watch?v=fiFCwc57MQs>
- [2] AirTrack https://www.youtube.com/watch?v=H3IL_Wjxjpw&t=1s
- [3] TrajAirNet <https://www.youtube.com/watch?v=eIAQXrxB2gw&t=15s>
- [4] Demo https://www.youtube.com/watch?v=iU_MyMwuE8E

5.4 Blogs:

- [1] How do you train AI Pilots? [\[Link\]](#)
- [2] XPlaneROS : ROS Wrapper for Autonomous Fixed Wing Applications [\[Link\]](#)
- [3] Long-range Aircraft Detection and Tracking [\[Link\]](#)
- [4] TrajAir: A General Aviation Trajectory Dataset [\[Link\]](#)

5.5 Media Coverage:

- [1] [Move over, autopilot: This AI can avoid other planes](#), Popular Science, Sept 2022
- [2] [CMU's AI pilot lands in the news](#), Practical AI – Episode #189, Sept 2022
- [3] [Researchers Develop AI Pilot for Navigating Crowded Airspace](#), Aviation Today, Aug 2022
- [4] [AI Pilot Can Navigate Crowded Airspace](#), CMU School Of Computer Science, Aug 2022

Certificate Of Completion

Envelope Id: D8FFA69230A84726BA12787788604873

Status: Completed

Subject: 2023-102 :: SIGNATURE REQUIRED :: CMU IP Disclosure

Source Envelope:

Document Pages: 20

Signatures: 8

Envelope Originator:

Certificate Pages: 6

Initials: 0

CMU CTTEC

AutoNav: Enabled

4615 Forbes Avenue, Ste 302

Enveloped Stamping: Enabled

4615 Forbes Avenue, Ste 302

Time Zone: (UTC-05:00) Eastern Time (US & Canada)

Pittsburgh, PA 15213

innovation@cmu.edu

IP Address: 108.39.251.101

Record Tracking

Status: Original

Holder: CMU CTTEC

Location: DocuSign

10/24/2022 8:44:26 AM

innovation@cmu.edu

Signer Events**Signature****Timestamp**

Ian Higgins



Sent: 10/24/2022 8:49:30 AM

ihiggins@andrew.cmu.edu

Resent: 10/31/2022 2:08:32 PM

Security Level: Email, Account Authentication

Resent: 11/10/2022 8:53:11 AM

(None)

Resent: 11/21/2022 3:05:34 PM

Signature Adoption: Drawn on Device
Using IP Address: 128.237.82.16

Viewed: 11/21/2022 3:10:22 PM

Signed: 11/21/2022 3:11:34 PM

Electronic Record and Signature Disclosure:

Accepted: 11/21/2022 3:10:22 PM

Sent: 10/24/2022 8:49:29 AM

ID: 464ef562-ad2c-41b2-972a-770002d13b4a

Ingrid Navarro Anaya



Viewed: 10/24/2022 6:54:47 PM

ingridn@andrew.cmu.edu

Security Level: Email, Account Authentication

Signed: 10/24/2022 6:55:59 PM

(None)

Signature Adoption: Pre-selected Style
Using IP Address: 14.10.146.128**Electronic Record and Signature Disclosure:**

Accepted: 10/24/2022 6:54:47 PM

Sent: 10/24/2022 8:49:29 AM

ID: 1060d52f-2fcf-4f41-85e2-d3032d3e91e8

Jasmine Jerry Aloor



Viewed: 10/24/2022 11:16:57 AM

jasminejerry@yahoo.in

Security Level: Email, Account Authentication

Signed: 10/24/2022 11:19:07 AM

(None)

Signature Adoption: Uploaded Signature Image
Using IP Address: 18.29.30.144**Electronic Record and Signature Disclosure:**

Accepted: 10/24/2022 11:16:57 AM

Sent: 10/24/2022 8:49:29 AM

ID: 91e49053-59e0-4f15-926e-54ad3954d63f

Jay Patrikar



Viewed: 10/24/2022 12:00:17 PM

jpatrika@andrew.cmu.edu

Security Level: Email, Account Authentication

Signed: 10/24/2022 12:01:04 PM

(None)


Signature Adoption: Pre-selected Style
Using IP Address: 71.182.232.115**Electronic Record and Signature Disclosure:**

Accepted: 10/24/2022 12:00:17 PM

Sent: 10/24/2022 8:49:29 AM

ID: f4b88d26-d777-446d-b3aa-46e41cf84722

Viewed: 10/24/2022 12:01:04 PM

Signer Events	Signature	Timestamp
Jean Oh hyaejino@andrew.cmu.edu Security Level: Email, Account Authentication (None)		Sent: 10/24/2022 8:49:28 AM Viewed: 10/24/2022 8:53:53 PM Signed: 10/24/2022 8:54:16 PM
Electronic Record and Signature Disclosure: Accepted: 10/24/2022 8:53:53 PM ID: a56d5c37-f7e7-40fa-9c4b-ce748d630072		
Joao Dantas jdantas@andrew.cmu.edu Security Level: Email, Account Authentication (None)		Sent: 10/24/2022 8:49:31 AM Viewed: 10/24/2022 9:34:35 AM Signed: 10/24/2022 9:35:44 AM
Electronic Record and Signature Disclosure: Accepted: 10/24/2022 9:34:35 AM ID: e3ceb150-0ca7-4677-8f99-d08bf535b8d3		
Rohan Bajjal rohanbajjal2702@gmail.com Security Level: Email, Account Authentication (None)		Sent: 10/24/2022 8:49:30 AM Resent: 10/31/2022 2:08:32 PM Resent: 11/10/2022 8:53:11 AM Viewed: 11/10/2022 9:23:25 AM Signed: 11/10/2022 9:25:43 AM
Electronic Record and Signature Disclosure: Accepted: 11/10/2022 9:23:25 AM ID: 8daf13df-9934-47c1-bfeb-4b4fafde5b93		
Sebastian Scherer basti@andrew.cmu.edu Security Level: Email, Account Authentication (None)		Sent: 10/24/2022 8:49:28 AM Viewed: 10/24/2022 9:00:27 AM Signed: 10/24/2022 9:00:33 AM
Electronic Record and Signature Disclosure: Accepted: 10/24/2022 9:00:27 AM ID: 559e414e-5fad-48f6-96ef-51b274041735		

In Person Signer Events	Signature	Timestamp
Editor Delivery Events	Status	Timestamp
Agent Delivery Events	Status	Timestamp
Intermediary Delivery Events	Status	Timestamp
Certified Delivery Events	Status	Timestamp
Carbon Copy Events	Status	Timestamp

Carbon Copy Events	Status	Timestamp
CMU CTTEC innovation@cmu.edu Center for Technology Transfer & Enterprise Creation Carnegie Mellon University Security Level: Email, Account Authentication (None)	COPIED	Sent: 10/24/2022 8:49:30 AM Resent: 11/21/2022 3:11:39 PM
Witness Events	Signature	Timestamp
Notary Events	Signature	Timestamp
Envelope Summary Events	Status	Timestamps
Envelope Sent Certified Delivered Signing Complete Completed	Hashed/Encrypted Security Checked Security Checked Security Checked	10/24/2022 8:49:31 AM 10/24/2022 9:00:27 AM 10/24/2022 9:00:33 AM 11/21/2022 3:11:34 PM
Payment Events	Status	Timestamps
Electronic Record and Signature Disclosure		

CONSUMER DISCLOSURE

From time to time, Carnegie Mellon University (we, us or Company) may be required by law to provide to you certain written notices or disclosures. Described below are the terms and conditions for providing to you such notices and disclosures electronically through the DocuSign, Inc. (DocuSign) electronic signing system. Please read the information below carefully and thoroughly, and if you can access this information electronically to your satisfaction and agree to these terms and conditions, please confirm your agreement by clicking the 'I agree' button at the bottom of this document.

Getting paper copies

At any time, you may request from us a paper copy of any record provided or made available electronically to you by us. You will have the ability to download and print documents we send to you through the DocuSign system during and immediately after signing session and, if you elect to create a DocuSign signer account, you may access them for a limited period of time (usually 30 days) after such documents are first sent to you. After such time, if you wish for us to send you paper copies of any such documents from our office to you, you will be charged a \$0.00 per-page fee. You may request delivery of such paper copies from us by following the procedure described below.

Withdrawing your consent

If you decide to receive notices and disclosures from us electronically, you may at any time change your mind and tell us that thereafter you want to receive required notices and disclosures only in paper format. How you must inform us of your decision to receive future notices and disclosure in paper format and withdraw your consent to receive notices and disclosures electronically is described below.

Consequences of changing your mind

If you elect to receive required notices and disclosures only in paper format, it will slow the speed at which we can complete certain steps in transactions with you and delivering services to you because we will need first to send the required notices or disclosures to you in paper format, and then wait until we receive back from you your acknowledgment of your receipt of such paper notices or disclosures. To indicate to us that you are changing your mind, you must withdraw your consent using the DocuSign 'Withdraw Consent' form on the signing page of a DocuSign envelope instead of signing it. This will indicate to us that you have withdrawn your consent to receive required notices and disclosures electronically from us and you will no longer be able to use the DocuSign system to receive required notices and consents electronically from us or to sign electronically documents from us.

All notices and disclosures will be sent to you electronically

Unless you tell us otherwise in accordance with the procedures described herein, we will provide electronically to you through the DocuSign system all required notices, disclosures, authorizations, acknowledgements, and other documents that are required to be provided or made available to you during the course of our relationship with you. To reduce the chance of you inadvertently not receiving any notice or disclosure, we prefer to provide all of the required notices and disclosures to you by the same method and to the same address that you have given us. Thus, you can receive all the disclosures and notices electronically or in paper format through the paper mail delivery system. If you do not agree with this process, please let us know as described below. Please also see the paragraph immediately above that describes the consequences of your electing not to receive delivery of the notices and disclosures electronically from us.

How to contact Carnegie Mellon University:

You may contact us to let us know of your changes as to how we may contact you electronically, to request paper copies of certain information from us, and to withdraw your prior consent to receive notices and disclosures electronically as follows:

To advise Carnegie Mellon University of your new e-mail address

To let us know of a change in your e-mail address where we should send notices and disclosures electronically to you, you must send an email message to us at innovation@cmu.edu and in the body of such request you must state: your previous e-mail address, your new e-mail address. We do not require any other information from you to change your email address..

In addition, you must notify DocuSign, Inc. to arrange for your new email address to be reflected in your DocuSign account by following the process for changing e-mail in the DocuSign system.

To request paper copies from Carnegie Mellon University

To request delivery from us of paper copies of the notices and disclosures previously provided by us to you electronically, you must send us an e-mail to innovation@cmu.edu and in the body of such request you must state your e-mail address, full name, US Postal address, and telephone number. Please provide your current email address so that we may personally send you documents for signature.

To withdraw your consent with Carnegie Mellon University

To inform us that you no longer want to receive future notices and disclosures in electronic format you may:

- i. decline to sign a document from within your DocuSign session, and on the subsequent page, select the check-box indicating you wish to withdraw your consent, or you may;
- ii. send us an e-mail to innovation@cmu.edu and in the body of such request you must state your e-mail, full name, US Postal Address, and telephone number. Please provide a reason why you are decline to consent to signature.. The consequences of your withdrawing consent for online documents will be that transactions may take a longer time to process..

Required hardware and software

Operating Systems:	Windows® 2000, Windows® XP, Windows Vista®; Mac OS® X
Browsers:	Final release versions of Internet Explorer® 6.0 or above (Windows only); Mozilla Firefox 2.0 or above (Windows and Mac); Safari™ 3.0 or above (Mac only)
PDF Reader:	Acrobat® or similar software may be required to view and print PDF files
Screen Resolution:	800 x 600 minimum

Enabled Security Settings:	Allow per session cookies
----------------------------	---------------------------

** These minimum requirements are subject to change. If these requirements change, you will be asked to re-accept the disclosure. Pre-release (e.g. beta) versions of operating systems and browsers are not supported.

Acknowledging your access and consent to receive materials electronically

To confirm to us that you can access this information electronically, which will be similar to other electronic notices and disclosures that we will provide to you, please verify that you were able to read this electronic disclosure and that you also were able to print on paper or electronically save this page for your future reference and access or that you were able to e-mail this disclosure and consent to an address where you will be able to print on paper or save it for your future reference and access. Further, if you consent to receiving notices and disclosures exclusively in electronic format on the terms and conditions described above, please let us know by clicking the 'I agree' button below.

By checking the 'I agree' box, I confirm that:

- I can access and read this Electronic CONSENT TO ELECTRONIC RECEIPT OF ELECTRONIC CONSUMER DISCLOSURES document; and
- I can print on paper the disclosure or save or send the disclosure to a place where I can print it, for future reference and access; and
- Until or unless I notify Carnegie Mellon University as described above, I consent to receive from exclusively through electronic means all notices, disclosures, authorizations, acknowledgements, and other documents that are required to be provided or made available to me by Carnegie Mellon University during the course of my relationship with you.