

# Language Conditioning Improves Accuracy of Aircraft Goal Prediction in Non-Towered Airspace

Sundhar Vinodh Sangeetha<sup>1</sup>, Chih-Yuan Chiu<sup>2</sup>, Sarah H.Q. Li<sup>1,†</sup>, Shreyas Kousik<sup>3,†</sup>

**Abstract**—Autonomous aircraft must safely operate in non-towered airspace, where coordination relies on voice-based communication among human pilots. Safe operation requires an aircraft to predict the intent, and corresponding goal location, of other aircraft. This paper introduces a multimodal framework for aircraft goal prediction that integrates natural language understanding with spatial reasoning to improve autonomous decision-making in such environments. We leverage automatic speech recognition and large language models to transcribe and interpret pilot radio calls, identify aircraft, and extract discrete intent labels. These intent labels are fused with observed trajectories to condition a temporal convolutional network and Gaussian mixture model for probabilistic goal prediction. Our method significantly reduces goal prediction error compared to baselines that rely solely on motion history, demonstrating that language-conditioned prediction increases prediction accuracy. Experiments on a real-world dataset from a non-towered airport validate the approach and highlight its potential to enable socially aware, language-conditioned robotic motion planning.

## I. INTRODUCTION

In recent years, there has been a surge of industry interest in operating autonomous aircraft from small, regional airports that operate without traffic control towers. These non-towered airports account for 92% of airports within the United States and 90% of airports worldwide [1]. Companies such as Reliable Robotics, Merlin Labs, and Ribbit have been actively developing systems for remotely piloted and fully autonomous aircraft, driven by the potential to reduce operational costs and expand access [2]. Despite the surging interest, safe integration of autonomous aircraft operations near non-towered airports remains incredibly challenging. This is in large part because non-towered airports lack towered airports' centralized surveillance, infrastructure, and air traffic controllers to ensure flight safety and resolve real-time flight path conflicts. Instead, potential conflicts are locally managed by pilots through a combination of visual flight rules and self-announced radio calls over the common traffic advisory frequency (CTAF), enabling pilots to detect and avoid one another. Therefore, flight safety near non-towered airports requires not only adherence to established aviation rules, but also the ability to interpret loosely structured natural language communications in real time.

All authors are with the Georgia Institute of Technology, Atlanta, GA, USA. <sup>1</sup> School of Aerospace Engineering. <sup>2</sup> School of Electrical and Computer Engineering. <sup>3</sup> School of Mechanical Engineering. † indicates equal advising. Corresponding author: ssangeetha3@gatech.edu.

**Codebase:**

<https://github.com/sundharvs/language-conditioned-aircraft-goal-prediction>

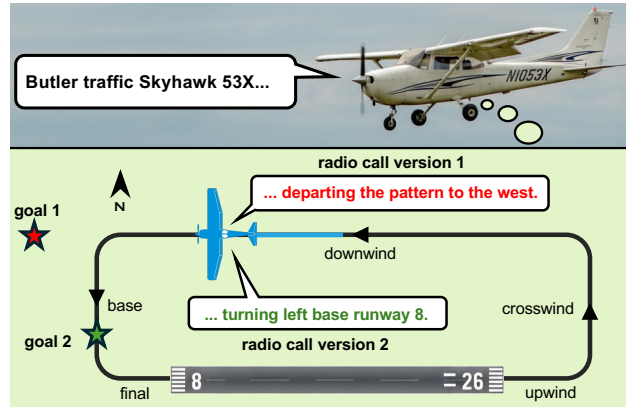


Fig. 1: Example of a pilot radio call on the CTAF in non-towered airspace. An aircraft (blue in inset) makes a radio call announcing its intent with respect to the “Runway 8” left traffic pattern. Multiple goals are possible depending on the radio call. In this work, we show that conditioning on such radio calls along with the aircraft’s recent trajectory (blue line on downwind leg) improves accuracy of goal prediction, which is critical for autonomous operation in non-towered airspace.

However, current state-of-the-art conflict avoidance approaches for autonomous aircraft strictly rely only on observed trajectories and structured air traffic rules [3], and do not consider radio communication between pilots, to the best of our knowledge. This means that, near non-towered airports, current methods will miss critical coordination information. This makes mixed human-autonomy operations dangerous, and potentially requires human pilots to accommodate and increase safety margins around autonomous aircraft. Towards solving this challenge, this paper takes key steps towards enabling robot aircraft to understand and act on human radio communication.

In this work, we aim to bridge the gap between unstructured human pilot communication and autonomous collision avoidance. That is, this work makes a key step towards addressing the broad question, how should an autonomous aircraft interpret and use radio communications to safely and effectively coordinate with other (potentially human-piloted) aircraft in a non-towered airport environment? This challenge can be broken into two research problems: predicting the motion of other aircraft using radio messages, and announcing the ego motion plan to other aircraft in a timely way while obeying radio communication rules. In this paper, we focus on the first problem by answering the following question in the affirmative: does incorporating language conditioning improve a robot aircraft’s ability to predict the motion of other aircraft in its vicinity?

### A. Related Work

Existing research in decentralized air traffic control primarily relies on historical observations to predict future trajectories [4]. Many works have recognized the importance of predicting intent to forecast future aircraft trajectories [5–17]. Generally, they approximate pilot intent as a function of the aircraft’s past trajectory. Foundational methods in this area explicitly model aircraft dynamics and use traditional statistical methods like Kalman filters and maximum likelihood inference to predict future trajectories [5]. Recent works have instead used deep learning approaches to build intent-driven prediction models with increasing complexity. For example, these models use additional information like maneuvering patterns from a repository of historical trajectories around an airport [6] and wind data [9]. They also employ advanced model architectures that have increased predictive power, such as encoder-decoder architecture [12], generative adversarial networks [13], and diffusion models [11]. Some recent works have recognized that aircraft radio communications contain intent information and developed trajectory prediction models conditioned on these communications [14–17]. However, these works consider direct instructions from air traffic controllers to aircraft, which map clearly to trajectory parameters like speed, altitude, and GPS waypoints. Instead, we consider non-towered airport communications, which are less structured.

We note that the robotics community at large is experiencing rapid advances in integrating natural language processing with perception and decision-making [18–20]. Recent developments in speech-to-text (STT) and large language models (LLMs) have enabled robots to interpret audio messages and leverage linguistic information for reasoning about their environment. Within this broader trend, our work seeks to demonstrate how natural language understanding can be effectively applied to the spatial trajectory prediction problem, thereby bridging the gap between unstructured human communication and structured motion planning.

### B. Contributions

We make three contributions:

- 1) an automatic speech recognition method to transcribe radio calls on CTAF and identify speakers,
- 2) a method to infer intent from radio calls in uncontrolled airspace and predict aircraft trajectories with reduced error, and
- 3) an experimental evaluation showing how language conditioning improves goal prediction performance on real-world non-towered airport data.

## II. PRELIMINARIES AND PROBLEM FORMULATION

Consider an autonomous aircraft operating in the airspace surrounding an uncontrolled airport, in which human-piloted aircraft are also operating. How can this aircraft forecast the future goals of other aircraft in this airspace using observations of their past trajectories as well as their natural language radio transmissions on the CTAF? We focus on predicting goals because it has been shown that accurate

goal prediction greatly improves the quality of downstream trajectory prediction for motion planning [11, 21–24].

Throughout this paper, we consider the prediction problem for a single aircraft, without conditioning on information about other aircraft. We note that other settings (e.g., autonomous cars) that do *not* have radio communication *do* rely on multi-agent conditioning [25]. However, since radio calls come from human pilots who are aware of each other socially, we assume that these calls implicitly encode potential interactions, and therefore we focus on predicting individual agents’ goals, leaving extensions to multi-agent prediction for future work.

To proceed, we first introduce our notation conventions, then formalize our problem statement, and finally briefly discuss details relevant to non-towered airspace.

### A. Notation

Let  $\mathbb{R}$  denote the reals and  $\mathbb{N}$  the natural numbers. Let  $\mathbb{L}$  denote the space of all possible natural language utterances. Consider an aircraft whose intent and goal we wish to predict. Let  $\mathbf{x}(t) = (x(t), y(t), z(t))$  denote its physical position at a time  $t \in \mathbb{R}$ . Let  $r(t) \in \mathbb{L}$  denote the natural language radio call made by the aircraft at time  $t$ . We use  $\mathcal{T} = \{t_0, \dots, t_{\text{obs}}\} \subset \mathbb{R}$  to denote the observation time horizon, such that we observe a trajectory  $(\mathbf{x}(t))_{t \in \mathcal{T}}$  of the aircraft<sup>1</sup>. Let  $t_r \leq t_{\text{obs}}$  denote the time of the aircraft’s latest radio call, and a time instance  $t_{\text{pred}} > t_{\text{obs}}$  to denote the chosen time step for goal prediction for all aircraft.

### B. Problem Statement and Solution Overview

The technical problem we seek to solve is how to predict an aircraft’s intent, represented as a goal position, given its recent trajectory and radio calls:

**Problem.** *Suppose we have observed an aircraft’s recent trajectory  $(\mathbf{x}(t))_{t \in \mathcal{T}}$  and latest radio call  $r(t_r) \in \mathbb{L}$ . We seek to predict the aircraft’s future goal  $\mathbf{g}$ , meaning its position at  $t_{\text{pred}} > t_{\text{obs}}$ , denoted*

$$\mathbf{g} = \mathbf{x}(t_{\text{pred}}). \quad (1)$$

To solve this problem, we model a probability distribution over the aircraft’s future goal position, conditioned on its observed trajectories and latest radio call, and approximate its  $\mathbf{g}$  as a sample  $\hat{\mathbf{g}}$  drawn from this distribution:

$$\mathbf{g} \approx \hat{\mathbf{g}} \sim p(\hat{\mathbf{g}} \mid (\mathbf{x}(t))_{t \in \mathcal{T}}, r(t_r)). \quad (2)$$

### C. Characterizing Non-Towered Terminal Airspace

As standardized by the FAA [26], for all non-towered airports in the US, the airspace surrounding the airport is structured around the traffic pattern. As shown in Figure 1, this terminal airspace is divided into upwind, crosswind, downwind, base, and final legs of the traffic pattern for each runway. Active runways also adopt left-handed or right-handed traffic conventions. For example, “left downwind” refers to the segment of the traffic pattern reached by making

<sup>1</sup>This information is publicly available from aircraft transponder data.

two left-hand 90° turns after departure, and “right downwind” refers to the segment reached by making two right-hand 90° turns after departure. Unless otherwise specified, we use the left-handed traffic pattern convention in the following sections, which is the default convention for non-towered airports.

Radio calls on CTAF frequencies in non-towered airspace typically contain four pieces of information: airport name, aircraft identifier, position, and intention. For example, “Butler County traffic Cherokee 135PL three miles south of the field entering left downwind runway 8.” Reported positions and intentions are typically a combination of distance, cardinal direction, and traffic pattern segment. Critically, pilots often abbreviate or omit some information, present it in various orders, and potentially insert verbal filler, all of which necessitates semantic understanding of radio calls.

### III. METHOD

In this section, we present our framework for predicting aircraft goals (i.e., trajectory endpoints) in non-towered airspace by conditioning on the observed past trajectories and natural language radio calls. Our approach consists of three components: (Sec. III-A) speech-to-text transcription and speaker aircraft identification to transform CTAF audio into text, (Sec. III-B) intent extraction using large language models with domain-specific context to reduce text into a compact intent representation, and (Sec. III-C) goal prediction with a learned, multimodal model, outputting a probability distribution of an aircraft goal.

#### A. Context Enhanced Speech-to-Text and Aircraft Identification

To incorporate pilot radio communications into aircraft goal prediction, it is essential to accurately transcribe audio messages and identify the transmitting aircraft. While trained speech-to-text ASR models exist for aviation radio communications [17, 27–29], we find that these models perform poorly on non-towered airspace radio calls, since they are trained on communications between air traffic controllers and commercial aircraft in controlled airspace [30]. Additionally, large general purpose speech-to-text models like Whisper [31] also do not perform well in transcribing radio calls due to audio noise and domain-specific terminology. However, we find that by providing large general purpose models with domain-specific information, we can accurately transcribe radio calls and identify transmitting aircraft.

Specifically, in our framework, a speech-to-text model and large language model are provided with *static context*  $s \in \mathbb{L}$  in the form of common terminology used in non-towered airspace, runway numbers, airport name, etc., plus 50 few-shot human-labeled examples of radio calls with corresponding intents. With each radio call, the models are also provided with *dynamic context*  $d \in \mathbb{L}$  in the form of a list of aircraft ADS-B identifiers (tail numbers) in the airspace surrounding the airport, their corresponding manufacturer and model names, which are retrieved from the FAA N-Number database, and their distance and cardinal direction

from the airport. This *dynamic context* is automatically generated at inference time. We define the total *domain context* as the combined static and dynamic contexts:

$$c = (s, d) \in \mathbb{L}. \quad (3)$$

We provide an example dynamic context in Appendix I; the static context is too large to print due to page limits but will be included in our open-source codebase.

The dynamic context is especially crucial in accurately identifying which aircraft transmitted a radio message. Commercial aircraft in controlled airspace must identify themselves in every radio call with the airline name and full flight number, which map directly to ADS-B identifiers. However, general aviation aircraft, which are the ones that largely operate at non-towered airports, have greater flexibility in how they identify themselves; using combinations of manufacturer names, model names and parts of their tail number. “N123AB”, “Cessna 23AB”, and “Skyhawk” are all examples of radio call identifiers for the same aircraft, whereas “N123AB” is the ADS-B identifier. Radio identification on CTAF is oriented towards visual identification of aircraft by human pilots, so identifiers like model names are most often used rather than full tail numbers.

#### B. Intent Modeling and Extraction

Because pilot radio communications are inherently unstructured, directly incorporating a radio call into trajectory prediction models greatly expands the input space and complicates training. To address this challenge, we transform radio calls into a compact form that captures the pilot’s intent, which we then use as an additional input to the goal prediction mechanism. This intent model aims to be more compact and uniform than loosely structured language sequences while remaining sufficiently rich to distinguish between different future spatial goals.

For a given aircraft, we consider its most recent radio call  $r(t_r) \in \mathbb{L}$  occurring at time  $t_r$ . If there has been no radio call for the past 10 minutes, we treat the radio call, and corresponding intent label as “unknown”.

We use a discrete set of intent labels to characterize the radio call  $r(t_r)$ . These labels specify distinct spatial approaches and locations with respect to the runway’s traffic pattern as shown in Fig. 1. Specifically, the intent labels captures (a) take off, landing, and entering three traffic pattern legs (crosswind, base, downwind) for each runway end (eg. runway 8 or 26 depending on the direction used), and (b) the four different cardinal directions of leaving the overall traffic pattern (e.g. NESW). Two additional labels, “insufficient information” and “other intent”, capture cases in which there is either insufficient information to determine intent or in which there is sufficient information, but the intent does not belong to one of the other classes. Empirically, we find that the “other intent” label predominantly captures ground operations that do not conflict with in-air operations.

Formally, let  $\mathcal{J} \subset \mathbb{L}$  denote the finite set of discrete intent labels. For each radio call  $r(t_r)$  from the aircraft, we extract the corresponding intent label by passing the radio

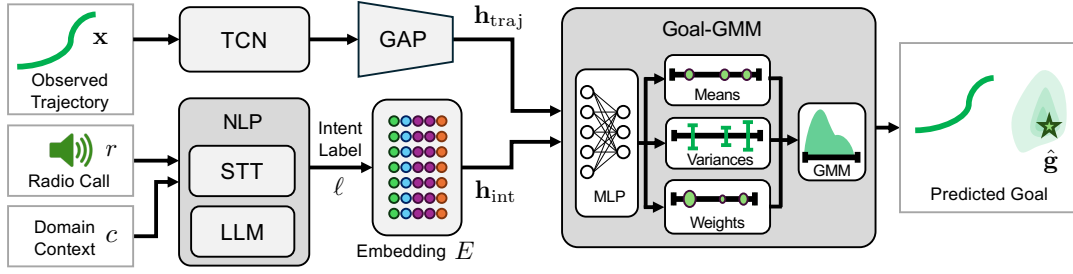


Fig. 2: An overview of our model architecture as described in Sec. III-C. Our key finding is that conditioning on natural language radio calls improves autonomous aircraft’s ability to predict the intent of other aircraft in non-towered airspace.

call through a speech-to-text (STT) model, then an LLM conditioned on the domain context:

$$\ell = \text{LLM} \left( \text{STT} (r(t_r)) \mid c \right) \in \mathcal{J}. \quad (4)$$

We use gpt-4o-transcribe for STT and Gemma 3 27B [32] as the LLM.

### C. Aircraft Goal Prediction

Given the observed trajectory of an aircraft and its language conditioned intent label  $\ell$ , our objective is to approximate the probability distribution as in (2) over the predicted future goal  $\hat{\mathbf{g}}$  of the aircraft. We approximate this distribution using a Gaussian Mixture Model (GMM), meaning a weighted sum of  $K$  independent Gaussian components, as

$$p(\hat{\mathbf{g}} \mid (\mathbf{x}(t))_{t \in \mathcal{T}}, r(t_r)) \approx \sum_{k=1}^K \pi_k \cdot \mathcal{N}(\mu_k, \Sigma_k \mid \mathbf{h}_{\text{traj}}, \mathbf{h}_{\text{int}}), \quad (5)$$

where we condition the distribution on latent vectors  $\mathbf{h}_{\text{traj}}$  and  $\mathbf{h}_{\text{int}}$  explained below. Each  $\pi_k \in [0, 1]$  is the mixture weight of the  $k^{\text{th}}$  component such that  $\sum_{k=1}^K \pi_k = 1$ . Each  $\mathcal{N}(\cdot)$  is a multivariate Gaussian distribution with mean  $\mu_k \in \mathbb{R}^3$  and covariance  $\Sigma_k \in \mathbb{R}^3$ . We assume a diagonal covariance matrix, such that  $\Sigma_k = \text{diag}((\sigma_k)^2)$  for a vector  $\sigma_k \in \mathbb{R}^3$ .

To implement (5), as shown in Figure 2, our proposed architecture processes the observed trajectory and intent label in parallel before combining them to predict the GMM parameters in the “GoalGMM” component. GoalGMM itself consists of three linear heads that output the GMM parameters: mean, variance, and mixture weight. Next, we detail the trajectory and intent processing and each prediction head.

**Trajectory Encoder:** To capture the motion patterns and temporal dependencies within the observed trajectory, we employ a Temporal Convolutional Network (TCN) [33]. The TCN processes each input trajectory  $(\mathbf{x}(t))_{t \in \mathcal{T}}$  using a stack of causal, dilated, 1-D convolutional layers with channel sizes of 64, 128, and 256. The TCN maps the input trajectory sequence to a sequence of high-dimensional feature vectors

$$(\mathbf{h}(t_0), \dots, \mathbf{h}(t_{\text{obs}})) = \text{TCN}((\mathbf{x}(t))_{t \in \mathcal{T}}), \quad (6)$$

where each  $\mathbf{h}(t) \in \mathbb{R}^{n_{\text{traj}}}$ , where  $n_{\text{traj}} = 256$  in our implementation. To produce a single, fixed-size representation, we apply global average pooling (GAP) over the time dimension:

$$\mathbf{h}_{\text{traj}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbf{h}(t). \quad (7)$$

The resulting vector  $\mathbf{h}_{\text{traj}}$  serves as a summary of the aircraft’s observed trajectory.

**Intent Embedding:** To incorporate intent information from radio calls, the discrete intent label  $\ell \in \mathcal{J}$  is mapped to a dense vector representation using a learned embedding layer. This allows the model to learn a meaningful, continuous representation for each of the  $|\mathcal{J}|$  possible intents. The embedding for the aircraft is produced as:

$$\mathbf{h}_{\text{int}} = E(\ell), \quad (8)$$

where  $E: \mathcal{J} \rightarrow \mathbb{R}^{n_{\text{int}}}$  is a learnable embedding dictionary and  $\mathbf{h}_{\text{int}} \in \mathbb{R}^{n_{\text{int}}}$  is the resulting intent vector. In our implementation,  $n_{\text{int}} = 64$ .

**Probabilistic Prediction Head:** These latent vectors  $\mathbf{h}_{\text{traj}}$  and  $\mathbf{h}_{\text{int}}$  are concatenated and passed through a shared Multi-Layer Perceptron (MLP) with hidden dimensions of [128, 64]. This representation is then fed in parallel to three distinct linear heads to predict the parameters of a GMM with  $K$  components:

- 1) *Mean Head:* Predicts the means  $\mu_k \in \mathbb{R}^{K \times 3}$  for each of the  $K$  Gaussian components.
- 2) *Variance Head:* Predicts variances as log-variances  $\log(\sigma^2)_k \in \mathbb{R}^{K \times 3}$  for numerical stability.
- 3) *Mixture Weight Head:* Predicts the logits for the mixing coefficients  $\pi_k \in [0, 1]$ , which determine the weight of each component.

In our implementation,  $K = 15$ . Putting these parts together, we have

$$\{\mu_k, \Sigma_k, \pi_k\}_{k=1}^K = \text{GoalGMM}(\mathbf{h}_{\text{traj}}, \mathbf{h}_{\text{int}}). \quad (9)$$

Then, the model is trained end-to-end by minimizing the negative log-likelihood loss  $\mathcal{L}$  of the ground-truth goal under the predicted GMM distribution, plus an entropy loss encouraging the GMM modes to diverge

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \log \left( \sum_{k=1}^K \pi_k \cdot f_{\mathcal{N}}(\mathbf{g}; \mu_k, \Sigma_k) \right) + \mathcal{L}_{\text{ent}} \quad (10)$$

where  $f_{\mathcal{N}}(x; \mu, \Sigma)$  denotes the PDF of a normal distribution with mean  $\mu$  and covariance  $\Sigma$  evaluated at  $x$ ,  $N$  is the number of goals per batch, and  $\mathcal{L}_{\text{ent}}$  is as in [34, Eq. (7)].

Training is performed for 150 epochs using the AdamW optimizer. The learning rate is initialized at  $1 \times 10^{-4}$  and reduced by a factor of 0.2 after 5 epochs without improvement in training loss.

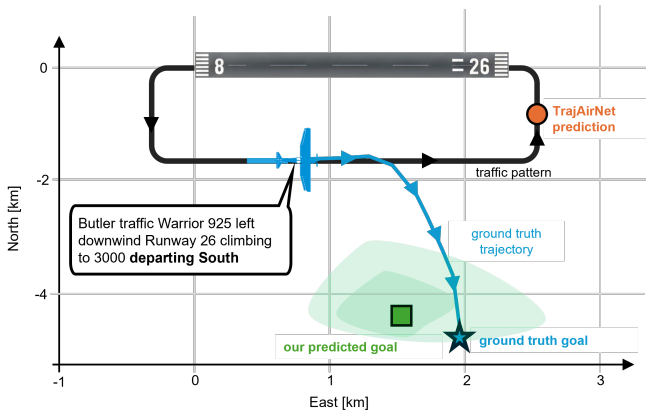


Fig. 3: Example trajectory prediction comparing our method to TrajAirNet [9], showing the improvement from using radio call info (in speech bubble).

#### IV. EXPERIMENTS

We now evaluate our method by testing speech recognition accuracy (Sec. IV-A), language-conditioned goal prediction (our core experiment, Sec. III-C), and feature/language ablations (Secs. IV-C and IV-D). All experiments are conducted on a workstation with a 32-core AMD Ryzen Threadripper PRO 7955WX CPU, 256 GB RAM, and an NVIDIA RTX 6000 Ada Generation GPU (48 GB VRAM).

##### A. Automatic Speech Recognition (ASR)

First, we evaluate the accuracy of our automatic speech recognition approach on transcribing historical radio communications on the CTAF observed at an uncontrolled airport, identifying the ADS-B IDs of aircraft that transmit messages, and labeling their intent from their messages.

*Hypothesis:* We hypothesize that given sufficient domain context<sup>2</sup>, general purpose ASR models can transcribe aviation radio communications well without any finetuning, and that large language models can identify which aircraft transmitted a radio call and accurately label its intent.

*Experiment Design:* We ran the experiments on a dataset of 250 radio call audios from the TartanAviation dataset [35] using gpt-4o-transcribe for speech-to-text and Gemma 3 27B for speaker identification and intent inference. The performance of our method is measured by comparing against human transcribed and annotated radio calls.

*Metrics:* We report the word error rate (WER) of our method, which is a common metric for ASR [36], the percentage of correctly identified speaker aircraft, and the percentage of correctly identified intent labels.

*Results and Discussion:* The results of this experiment are summarized in Table I. We observe that including domain context through gpt-4o-transcribe and Gemma 3 27B prompt input improves transcription, speaker identification, and intent labeling accuracy. It has been noted in the automatic speech recognition literature that minor differences in transcripts can lead to large word error rates despite accurately capturing meaning [31]. Our results support this notion and

<sup>2</sup>Domain context is defined in Sec. III-A. An example is provided in Appendix I

Method	WER	SIA (%)	ILA (%)
With domain context	<b>33.97</b>	<b>94.8</b>	<b>82.8</b>
Without domain context	60.55	63.6	32.8

TABLE I: ASR accuracy of aviation radio communications in uncontrolled airspace. SIA denotes speaker identification accuracy, ILA denotes intent labeling accuracy. Including domain-specific context significantly improves speech-to-text, speaker identification, and intent labeling performance.

suggest that these errors do not have a large impact on the ability of LLMs to accurately extract meaning from text. Despite relatively large word error rates, we achieve high accuracy in speaker identification and intent labeling from transcription.

##### B. Language Conditioned Goal Prediction

In this experiment, we test our goal prediction model conditioned on aircraft intent labels extracted from transcribed CTAF radio communications. We feed real-world aircraft trajectory data and radio call transcripts into our prediction model and compare its outputs against ground-truth trajectories.

*Hypothesis:* We hypothesize that conditioning the aircraft goal prediction model on intent labels from radio calls will decrease final displacement error (FDE) of predicted future goals compared to only conditioning the model on the position histories and wind contexts of each aircraft.

*Experiment Design:* We tested on aircraft trajectory and audio data collected at the Pittsburgh-Butler Regional Airport (KBTP) from the TartanAviation dataset [35]. A number of works [9–11] test their methods on a subset of this data made available in an earlier publication, the TrajAir dataset [9], labeled as “7days-1”, “7days-2”, “7days-3”, and “7days-4”. However, audio files for many time periods, including entire days, are missing from this subset. Importantly, “7days-4” does not have any audio data and “7days-1” only has one day of audio data, so we do not report results for these subsets. To compare our method fairly to these previous works, we evaluate on the subset containing audio. Additionally, we conduct our other experiments on a one week subset from the larger TartanAviation dataset (labeled 7daysJune) that contains complete audio and trajectory data.

As in TrajAirNet [9], the observation time horizon is set to 11 seconds, and the prediction time horizon is set to 120 seconds. We study the impact of changing these parameters in Sec. IV-D.

*Metrics:* We use a standard metric for trajectory prediction to quantify the performance of our model: best-of- $N$  (with  $N = 10$ ) final displacement error (FDE) [37], defined as the average over all aircraft of the lowest error,

$$\text{FDE} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \min_{j \in \{1, \dots, N\}} \|\mathbf{g}_{a,j} - \hat{\mathbf{g}}_{a,j}\|_2, \quad (11)$$

where  $\mathcal{A} = \{1, \dots, n_{\text{agt}}\}$  indexes all aircraft in the airspace. This metric is chosen to allow comparison of trajectory prediction performance to a large number of existing methods

by averaging across all goal positions predicted across an entire test dataset.

However, FDE alone does not capture the variance in the prediction performance across different testing sets, which is critical to understand a method’s precision for downstream motion planning; that is, if a method has high variance, then it is more difficult to use for planning collision-avoiding motions, even if it is accurate. Thus, as measures of dispersion in FDE, we report standard deviation in tables and show inter-quartile range (IQR) (to visualize skew) in figures.

*Baselines:* The main baseline we compare to is TrajAirNet [9], which has open source code; that work was one of the first to consider this specific problem, and introduced constant velocity and nearest neighbor baselines. Following the literature, we also report results claimed by several baselines directly from the corresponding papers [7–11], as code was either not available or could not be run without errors; we are working on reproducing these methods’ results for future work. We further note that these do not report any variance or dispersion measures in their FDE, which makes it difficult to understand or compare against how precise these methods are. To handle this, we rerun TrajAirNet to confirm their results and compute variance.

Algorithm	7Days-1	7Days-2	7Days-3	7Days-4
Const. Vel. [25] c.f. [9]	4.08	4.31	4.30	4.16
Nearest Neigh. [9]	2.70	1.99	2.69	2.58
STG-CNN [7] c.f. [9]	2.35	2.70	2.67	2.29
TransformerTF [8] c.f. [9]	3.85	4.10	4.36	4.19
TrajAirNet [9]	1.42	1.63	1.72	1.41
Social-PatteRNN-ATT [10]	1.42	1.67	1.65	1.51
GoodFlight [11]	0.41	0.40	0.48	0.40
TrajAirNet [9] (rerun)	1.46 ± 0.63	1.72 ± 0.60	1.91 ± 0.61	1.64 ± 0.77
<b>Ours</b>	-	0.71 ± 0.81	0.69 ± 0.57	-

TABLE II: Goal prediction error (FDE) [km].

Method	Mean FDE	25th pctl.	75th pctl.	St. Dev.
TrajAirNet [9]	1.390	0.860	1.731	0.686
<b>Ours</b>	<b>0.486</b>	<b>0.160</b>	<b>0.578</b>	<b>0.453</b>

TABLE III: Goal prediction error (FDE) [km] in the 7daysJune subset.

*Results and Discussion:* The results of this experiment are summarized in Table II. Across all datasets, we show significant improvement over all but one of the baselines, GoodFlight [11], and note that the improvement in FDE falls outside the standard deviation, suggesting a statistically significant improvement.

Although we do not achieve higher performance than in GoodFlight, we note that our method to infer goals from language is complementary to the goal-conditioned model in GoodFlight, and that our primary goal is to show the improvement in goal prediction performance as a result of language conditioning. We also note that the transformer-based diffusion model in GoodFlight has a much larger model capacity, increasing its predictive power and perhaps yielding an even larger improvement if combined with our language conditioning. However, as the code for this model is not open-source, we are unable to test this hypothesis.

Finally, we test our model on the 7daysJune dataset with complete audio and trajectory data and report lower error as well as dispersion statistics in Table III. An example goal prediction is shown in Fig. 3.

### C. Feature Importance Study

We evaluate the impact of language conditioning on goal prediction by conducting a feature importance study on the intent labels extracted from radio calls, which serve as features for our model. Although it is difficult to map exactly how specific inputs affect model output with a deep learning model, we evaluate feature importance using two well-known methods: Permutation Feature Importance (PFI) [38, §10] and Leave One Feature Out (LOFO).

*Hypothesis:* We hypothesize that removing intent information degrades goal prediction performance.

*Experiment Design:* In practice, evaluating PFI involves testing the pre-trained multimodal model on a dataset with shuffled aircraft intent labels. The corresponding increase in prediction error is used to characterize the reliance of the model on intent label features.

Evaluating LOFO involves training a new model that is not conditioned on intent labels and only on aircrafts’ past trajectories. The difference in error relative to the full model characterizes importance.

*Metrics:* For evaluating PFI and LOFO, we track change in the previously defined prediction error metric: FDE.

*Results and Discussion:* Ablating away the language-conditioned intent labels shows the impact of this information in improving goal prediction performance. With shuffled intent labels (PFI), we observe a modest increase in best-of- $N$  FDE of +0.163 compared to the baseline. This indicates that our model makes systematic use of intent features.

With intent features entirely removed (LOFO), the degradation is much larger: sampled FDE increases by +0.598 on average (95% CI [0.591, 0.604],  $p < 0.001$ ). This result confirms that intent labels provide substantial predictive value beyond trajectory information alone.

### D. Performance Under Varying Information Conditions

We observe that many previous works in trajectory prediction only present results where certain specific observation and prediction time horizons are used (c.f., [9–11]). Instead, we aim to study the change in our model’s performance as information parameters such as these are changed. We analyze the performance of our model by evaluating its sensitivity to several key factors:

- 1) The observation time horizon,  $t_{\text{obs}}$ .
- 2) The prediction time horizon,  $t_{\text{pred}}$ .
- 3) The time difference between the prediction time and the time the most recent radio call was received.
- 4) Radio calls with “Unknown” extracted intent labels.

*Hypothesis:* We hypothesize that our language-conditioned trajectory prediction model architecture is robust to changes in observation time horizon, prediction time horizon, and time difference between inference and radio call, and that trajectory prediction error is larger for “Unknown” intent labels.

*Results and Discussion:* As shown in Fig. 4, the goal prediction error (FDE) of our model does not change significantly with an increased observation time horizon. This result suggests that future goals are not correlated with the past position history, but rather with the current position and future intent. In Fig. 5, however, goal prediction performance does change with the length of the prediction horizon. While both the language conditioned model and trajectory-only model experience loss in accuracy when trying to predict longer horizon goals, this loss is greater for the trajectory-only model.

Finally, the impact of the recency of radio calls is shown in Fig. 6. Since radio calls are less frequently received than ADS-B position data, we condition predictions on radio calls that are up to 10 minutes old. However, we see that using older radio calls does not result in significantly worse prediction performance, supporting the idea that radio calls contain long horizon intentions which persist over time.

However, for radio calls with “Unknown” intent labels, our model predicts goal positions with an FDE of  $0.5832 \pm 0.6544$  km whereas for all other intent labels the FDE is  $0.4836 \pm 0.5131$  km. This result suggests that the prediction performance is worse without intent information.




Fig. 4: Final displacement error (FDE) of predicted aircraft goals with respect to observation time horizon. Longer observation windows do not improve goal prediction accuracy. The observation horizons used in TrajAirNet [9], GooDFlight [11] and Social-PatTeRNN-ATT [10] are shown as vertical lines. IQR is plotted as a measure of dispersion.

## V. CONCLUSION

The key finding of this paper is that conditioning on natural language radio calls can improve predictions of aircraft desired goal locations. To achieve this, we proposed a new deep neural network model to approximate a distribution over possible goal positions based on an aircraft’s recent trajectory history, its most recent radio call, and domain context. Experiments showed that such conditioning on language improves goal prediction accuracy over baselines; we further confirm that ablating away language significantly decreases accuracy, and language conditioning improves prediction accuracy further into the future. This improved prediction accuracy is a




Fig. 5: Final displacement error (FDE) of predicted aircraft goals vs. prediction time horizon length. Error increases with longer horizons, highlighting the challenge of long-term forecasting. However, language-conditioned goal prediction is better suited to address this challenge, showing a smaller increase in error. The prediction horizons used in TrajAirNet [9], GooDFlight [11] and Social-PatTeRNN-ATT [10] are shown as vertical lines. IQR is plotted as a measure of dispersion.

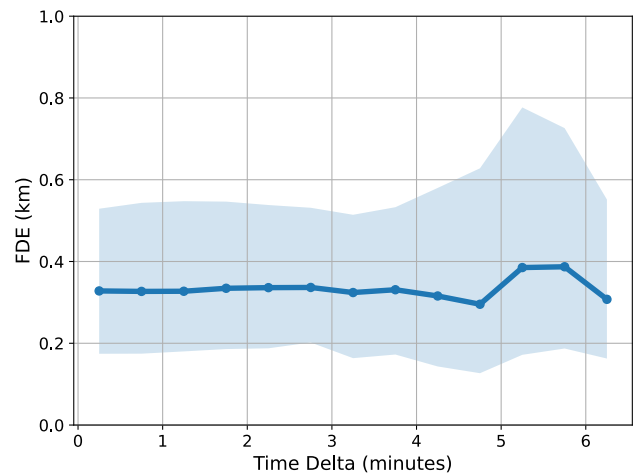


Fig. 6: Final displacement error (FDE) of predicted aircraft goals vs. the time difference between the most recent radio call used for prediction and the time at which prediction is made. Our model does not exhibit significantly increased error with longer delay. IQR is plotted as a measure of dispersion.

critical step towards satisfying proposed safety standards for autonomous aircraft, which stipulate a minimum separation distance of 1,500 feet (0.457 km) in terminal airspace [39].

Future work will address several key limitations. First, we have not integrated goal prediction into a closed-loop autonomous pipeline. Furthermore, our approach is specific to a single airport (which is a challenge shared by all of the baselines). All baselines also currently share large variance in their predictions. Finally, it remains open how best to autonomously generate correct, informative radio calls.

## REFERENCES

- [1] Federal Aviation Administration, “Air Traffic by the Numbers: ATO Fact Book — FY 2023,” U.S. Department of Transportation, Federal

- Aviation Administration, Air Traffic Organization, ATO Fact Book (Tech. Report), May 2024, 8th edition.
- [2] J. C.-H. Wang et al., "Flight testing an optionally piloted aircraft: A case study on trust dynamics in human-autonomy teaming," *arXiv preprint arXiv:2503.16227*, 2025.
  - [3] X. Yang and P. Wei, "Autonomous free flight operations in urban air mobility with computational guidance and collision avoidance," *IEEE Transactions on Intelligent transportation systems*, vol. 22, no. 9, pp. 5962–5975, 2021.
  - [4] P. Park, H. Khadilkar, H. Balakrishnan, and C. Tomlin, "Hybrid communication protocols and control algorithms for nextgen aircraft arrivals," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 2, pp. 615–626, 2014.
  - [5] J. L. Yepes, I. Hwang, and M. Rotea, "New algorithms for aircraft intent inference and trajectory prediction," *Journal of guidance, control, and dynamics*, vol. 30, no. 2, pp. 370–382, 2007.
  - [6] Y. Yin, S. Zhang, Y. Zhang, Y. Zhang, and S. Xiang, "Aircraft trajectory prediction in terminal airspace with intentions derived from local history," *Neurocomputing*, vol. 615, p. 128 843, 2025.
  - [7] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Socialstgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14424–14432.
  - [8] F. Giuliari, I. Hasan, M. Cristani, and F. Galasso, "Transformer networks for trajectory forecasting," in *2020 25th international conference on pattern recognition (ICPR)*, IEEE, 2021, pp. 10335–10342.
  - [9] 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," in *2022 international conference on robotics and automation (icra)*, IEEE, 2022, pp. 2525–2531.
  - [10] I. Navarro and J. Oh, "Social-patternn: Socially-aware trajectory prediction guided by motion patterns," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2022, pp. 9859–9864.
  - [11] S. Yang, L. Liu, B. Chen, S. Cheng, Z. Shi, and Z. Zou, "Goodflight: Goal-oriented diffusion model for flight trajectory prediction," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.
  - [12] P. N. Tran, H. Q. Nguyen, D.-T. Pham, and S. Alam, "Aircraft trajectory prediction with enriched intent using encoder-decoder architecture," *IEEE Access*, vol. 10, pp. 17 881–17 896, 2022.
  - [13] Y. Pang and Y. Liu, "Conditional generative adversarial networks (cgan) for aircraft trajectory prediction considering weather effects," *AIAA Scitech 2020 Forum*, 2020.
  - [14] Y. Wang et al., "A voice communication-augmented simulation framework for aircraft trajectory simulation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 7310–7320, 2021.
  - [15] Y. Pang, A. P. Kendall, A. Porcayo, M. Barsotti, A. Jain, and J.-P. Clarke, "From voice to safety: Language ai powered pilot-atc communication understanding for airport surface movement collision risk assessment," *arXiv preprint arXiv:2503.04974*, 2025.
  - [16] D. Guo, Z. Zhang, B. Yang, J. Zhang, H. Yang, and Y. Lin, "Integrating spoken instructions into flight trajectory prediction to optimize automation in air traffic control," *Nature Communications*, vol. 15, no. 1, p. 9662, 2024.
  - [17] J. Rakas, V. K. Vallioor, J. Krozel, P. F. Kostiuk, and M. T. Mohen, "Controller-pilot voice communication and intent monitoring for future aviation systems safety," in *AIAA AVIATION FORUM AND ASCEND 2024*, 2024, p. 3942.
  - [18] J. Wang et al., "Large language models for robotics: Opportunities, challenges, and perspectives," *Journal of Automation and Intelligence*, vol. 4, no. 1, pp. 52–64, 2025.
  - [19] M. J. Kim et al., "Openvla: An open-source vision-language-action model," in *Proceedings of The 8th Conference on Robot Learning*, P. Agrawal, O. Kroemer, and W. Burgard, Eds., ser. Proceedings of Machine Learning Research, vol. 270, PMLR, Nov. 2025, pp. 2679–2713.
  - [20] Y. Kim, D. Kim, J. Choi, J. Park, N. Oh, and D. Park, "A survey on integration of large language models with intelligent robots," *Intelligent Service Robotics*, vol. 17, no. 5, pp. 1091–1107, 2024.
  - [21] K. Mangalam et al., "It is not the journey but the destination: Endpoint conditioned trajectory prediction," in *European conference on computer vision*, Springer, 2020, pp. 759–776.
  - [22] C. Choi, S. Malla, A. Patil, and J. H. Choi, "Drogon: A trajectory prediction model based on intention-conditioned behavior reasoning," in *Proceedings of the 2020 Conference on Robot Learning*, J. Kober, F. Ramos, and C. Tomlin, Eds., ser. Proceedings of Machine Learning Research, vol. 155, PMLR, Nov. 2021, pp. 49–63.
  - [23] D. Li, Q. Zhang, S. Lu, Y. Pan, and D. Zhao, "Conditional goal-oriented trajectory prediction for interacting vehicles," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
  - [24] A. Fadillah, C.-L. Lee, Z.-X. Wang, and K.-T. Lai, "Goalnet: Goal areas oriented pedestrian trajectory prediction," *IEEE Access*, vol. 13, pp. 132 537–132 546, 2025.
  - [25] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conference on Computer Vision*, Springer, 2020, pp. 683–700.
  - [26] Federal Aviation Administration, "Non-towered airport flight operations," U.S. Department of Transportation, Federal Aviation Administration, Tech. Rep. AC 90-66B, Feb. 2018, Advisory Circular.
  - [27] S. Badrinath and H. Balakrishnan, "Automatic speech recognition for air traffic control communications," *Transportation research record*, vol. 2676, no. 1, pp. 798–810, 2022.
  - [28] S. V. Subramanian, P. F. Kostiuk, and G. Katz, "Custom ibm watson speech-to-text model for anomaly detection using atc-pilot voice communication," in *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 3979.
  - [29] J.L.P.M. van Doorn, *Whisper-large-v3-atco2-asr*, 2023.
  - [30] J. Zuluaga-Gomez, K. Vesely, I. Szöke, P. Motlicek, et al., "Atco2 corpus: A large-scale dataset for research on automatic speech recognition and natural language understanding of air traffic control communications," *arXiv preprint arXiv:2211.04054*, 2022.
  - [31] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning*, PMLR, 2023, pp. 28 492–28 518.
  - [32] G. Team et al., "Gemma: Open models based on gemini research and technology," *arXiv preprint arXiv:2403.08295*, 2024.
  - [33] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," in *Computer Vision – ECCV 2016 Workshops*, G. Hua and H. Jégou, Eds., Cham: Springer International Publishing, 2016, pp. 47–54.
  - [34] Y. Zhou, J. Gao, and T. Asfour, "Movement primitive learning and generalization: Using mixture density networks," *IEEE Robotics & Automation Magazine*, vol. 27, no. 2, pp. 22–32, 2020.
  - [35] J. Patrikar et al., "Image, speech, and ads-b trajectory datasets for terminal airspace operations," *Scientific Data*, vol. 12, p. 468, 2025.
  - [36] J. G. Fiscus, "A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover)," in *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, IEEE, 1997, pp. 347–354.
  - [37] D. Zhao and J. Oh, "Noticing motion patterns: A temporal cnn with a novel convolution operator for human trajectory prediction," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 628–634, 2020.
  - [38] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
  - [39] RTCA Special Committee 228, "DO-365: Minimum Operational Performance Standards (MOPS) for Detect and Avoid (DAA) Systems," RTCA, Inc., DO-365, 2017.

## APPENDIX I

Example of dynamic context, the time-varying part of the domain context  $d$ , used to identify the transmitting aircraft from a radio call:

Your task is to identify which aircraft made the call from a set of options. Only reply with the aircraft tail number with no additional text, if you cannot make a determination return Unknown. Options: N17NA - P28A, Piper, Carioquinha, Cherokee, Liner, 140-4, Chief, Archer, Cruiser, Challenger, Cadet, Warrior, Flite, Tupi Location - 5 miles, North; N2423U - C172, Cutlass, Hawk, Rocket, Skyhawk, Mescalero, Reims, Powermatic Location - 3 miles, East