



OPEN

DATA DESCRIPTOR

Image, speech, and ADS-B trajectory datasets for terminal airspace operations

Jay Patrikar¹✉, Joao Dantas², Brady Moon¹, Milad Hamidi¹, Sourish Ghosh¹, Nikhil Keetha¹, Ian Higgins¹, Atharva Chandak¹, Takashi Yoneyama² & Sebastian Scherer¹

We introduce TartanAviation, an open-source multi-modal dataset focused on terminal-area airspace operations. TartanAviation provides a holistic view of the airport environment by concurrently collecting image, speech, and ADS-B trajectory data using setups installed inside airport boundaries. The datasets were collected at both towered and non-towered airfields across multiple months to capture diversity in aircraft operations, seasons, aircraft types, and weather conditions. In total, TartanAviation provides 3.1M images, 3374 hours of Air Traffic Control speech data, and 661 days of ADS-B trajectory data. In addition to the raw data, we provide post-processed versions with synchronized, filtered, and interpolated data. In addition to the dataset, we also open-source the code-base used to collect and pre-process the dataset, further enhancing accessibility and usability. We believe this dataset has many potential use cases and would be particularly vital in allowing AI and machine learning technologies to be integrated into air traffic control systems and advance the adoption of autonomous aircraft in the airspace.

Background & Summary

In 2023, an average of 45,000 flights per day served over 2.9 million passengers in the US¹, which is nearly at the saturation capacity of the current US air traffic control system². The expected introduction of Advanced Aerial Mobility (AAM) operations within the National Airspace System heralds a further increase in flight operations, necessitating a need to increase the airspace capacity³. Compared to en-route airspace, terminal areas experience a higher air traffic density as nearly all aerial operations start or end within these areas. The effective future management of this high-risk airspace necessitates insights into various aspects of air traffic management. In addition to manned AAM operation, data is needed to achieve advances in fully or partially autonomous AAM agents. Curated datasets within this domain not only enable analytical research into understanding how the system is currently managed but also help accelerate the development of novel technologies for future operations. More generally, the aviation domain presents newer challenges to widely applied technologies like vision-based object detection, speech-to-text translation, and time-series analytics. Advancements in AI and machine learning can potentially revolutionize air traffic control systems, ensuring safer and more efficient coordination for the ever-increasing number of flights.

In this work, we introduce TartanAviation, a multi-modal dataset collected at towered and non-towered terminal areas within the US. The TartanAviation dataset covers three primary concurrently collected modalities of data: trajectory positions for capturing the spatial and temporal information of aircraft movements, video flight sequences collected with static cameras installed within terminal areas, and audio communications to document the voice interactions between pilots and air traffic controllers. While prior datasets in the aviation domain have focused on specific modalities like speech^{4,5} or vision^{6,7}, TartanAviation aims to provide a more holistic view of terminal airspace operations across various data modalities. Additionally, while previous datasets focus on large commercial airports, TartanAviation focuses on smaller regional airports within the Greater Pittsburgh area. Regional airports serve a multitude of different aircraft and mission profiles, providing a richer and more diverse data stream. We specifically focus on two airports: the towered Allegheny County Airport (ICAO:KAGC) and the non-towered Pittsburgh-Butler Regional Airport (ICAO:KBTP).

¹Robotics Institute, Carnegie Mellon University, Pennsylvania, USA. ²Aeronautics Institute of Technology, Sao Jose dos Campos, SP, Brazil. ✉e-mail: patrikarjay@gmail.com

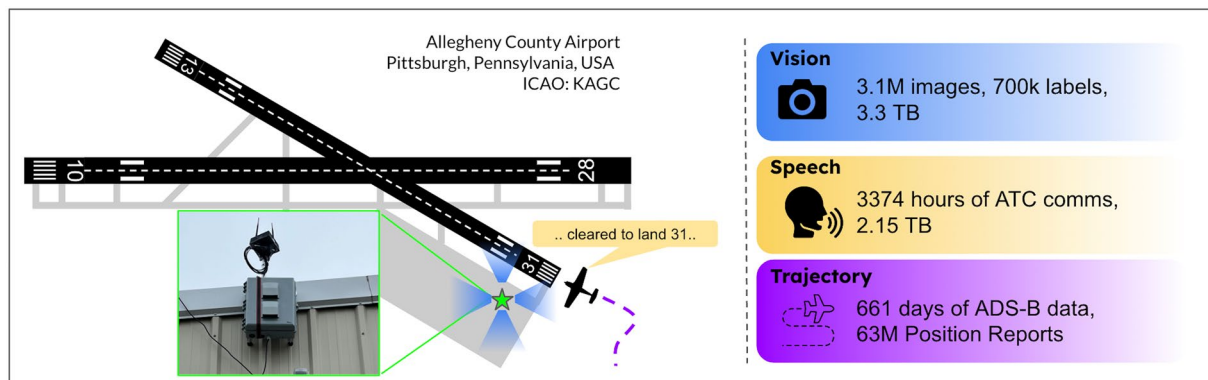


Fig. 1 Our custom data collection setup installed at the Allegheny County Airport with its approximate location within the airport premises with respect to the runway geometry. The setup recorded images, audio, and aircraft trajectory data.

Human vision is the last line of defense against mid-air collisions, making it critical for aviation safety⁸. With recent advances in computer vision technologies, visual detect-and-avoid (DAA) systems have shown promising results in detecting airborne objects at greater distances⁹. Prior datasets for DAA either lack diversity⁷, are computer-generated⁶, or are relatively small^{10,11}.

In this context, the TartanAviation vision dataset is a large-scale real-world dataset collected by placing static cameras within the terminal area. TartanAviation currently offers 3.1 million images covering challenging scenarios like snow, mist, rain, varying cloud cover, and diverse aircraft types. The data is augmented with the associated weather and air traffic trajectory data. With over 700k aircraft labels, these challenging real-world scenarios promise to enable more research in robust computer vision techniques for long-range object detection.

Trajectory data represents the time-series information of aircraft positions operating within the terminal airspace. This data is collected by recording Automatic Dependent Surveillance-Broadcasts (ADS-B) using receivers placed within the terminal area. ADS-B is an aviation surveillance technology where aircraft automatically broadcast their position and other data using GNSS satellite navigation, enabling tracking by air traffic control and nearby aircraft without ground interrogation. It is a key component of global air traffic systems and is mandatory in several countries for enhanced safety and efficiency. Prior studies have looked into the performance of ADS-B including its accuracy, integrity, continuity and availability¹² as well as its use in sequencing the optimal landing order for arriving flights.¹³ Our prior work, TrajAir¹⁴, provided 111 days of ADS-B data at KBTP. TartanAviation extends this tally to a total of 661 days of data at both KBTP and KAGC. Beyond the aviation domain, trajectory data can enable more research in time-series forecasting, social trajectory prediction, and anomaly detection.

Radio communications enable pilots and Air Traffic Controllers (ATC) to share time-sensitive intent and instructions over dedicated frequencies. The information included in these communications helps pilots and ATC to build situational awareness, enabling the safe operation of aircraft on the ground and in the air. Radio communication within aviation use both High-Frequency (HF) as well as Very-High-Frequency (VHF) bands. While HF is mostly used for long-range communication, for example in oceanic flights, VHF is preferred for short to medium range communication. A VHF radio operates on frequencies between 118.0 megahertz (MHz) and 136.975 MHz¹⁵. While most speech datasets are unimodal¹⁶, recent multi-modal datasets with text and trajectory¹⁷ have focused exclusively on commercial aircraft operations at large airports. TartanAviation provides first-of-its-kind speech data at relatively smaller airports, including both towered and untowered airfields. Along with the diversity in speech data, TartanAviation also provides concurrent trajectory data, enabling novel research in multi-modal speech-to-text translation and speech-to-intent predictions¹⁸ conditioned on external context.

All of the data in TartanAviation was collected with administrative support and prior authorization at both the KAGC and KBTP airports. TartanAviation represents our first step at building a centralized comprehensive multi-modal repository for aircraft data within the terminal airspace.

Method

This section briefly overviews the equipment and protocols used to collect the datasets. Figure 2 shows the setup and an overview of the data collection pipeline. Prior to data collection, the Carnegie Mellon University Institutional Review Board (IRB) determined that this project does not involve human subjects research, as defined by federal regulations. In the United States, ADS-B transmissions and ATC communications are publicly accessible broadcasts. According to 18 U.S.C. § 2511(2)(g)(ii)(IV)¹⁹, it is lawful to intercept radio communications transmitted by any marine or aeronautical communications system. Similar rules apply for ADS-B transmissions²⁰. Therefore, recording and sharing ADS-B data and ATC communications are permissible under U.S. law. Researchers in other jurisdictions should consult local laws and regulations, as the legality of recording and sharing such data may vary.

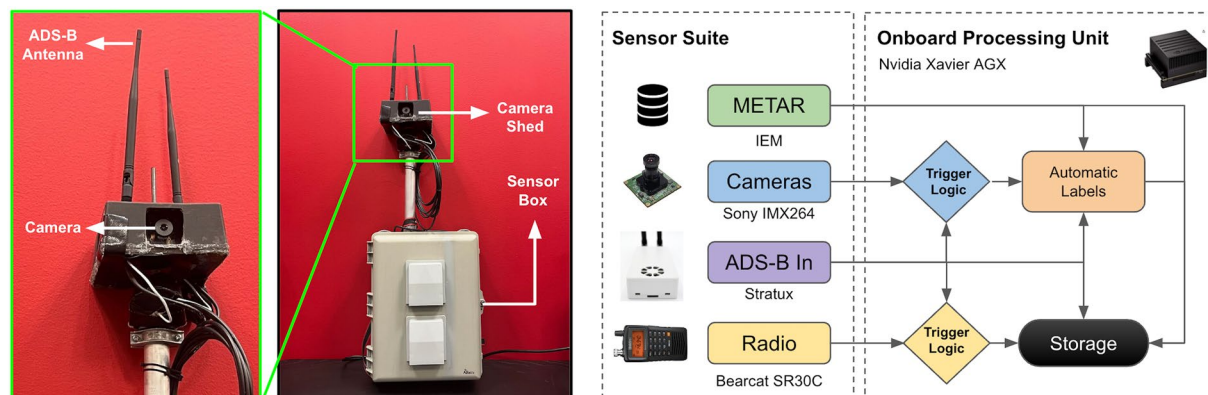


Fig. 2 Our custom setup hardware with the camera and ADS-B antenna mounts (*left*). We also showcase the data collection pipeline with the associated sensor suite and automatic logic that triggers camera and speech recordings (*right*).

Extension	Nomenclature	Content
.zip	<camera_id>_<timestamp>	Zip folder containing sequence data
↔ .mp4	<camera_id>_<timestamp>	Video File
↔ .avi	<camera_id>_<timestamp>_sink_verified	Video File with embedded labels
↔ .srt	<camera_id>_<timestamp>_subtitle	Raw timestamps of the recorded video
↔ .pkl	<camera_id>_<timestamp>_sink_adsb	Raw ADS-B dictionary
↔ .pkl	<camera_id>_<timestamp>_acft_sink	Raw aircraft type data
↔ .zip	<camera_id>_<timestamp>_labels	Zip folder containing the image labels
↔ .label	<frame_number>.label	Text file containing label data

Table 1. File structure information for TartanAviation image data.

Vision Data Acquisition. The vision data was exclusively collected at the KAGC airport. The setup was mounted at 40.351422 latitude, -79.923939 longitude, as shown in Fig. 1. The vision setup uses an array of Sony IMX 264 cameras connected to a NVIDIA Xavier AGX 32GB dev kit via a LI-JXAV-MIPI-ADPT-6CAM camera adaptor. The cameras collect data at 2048×2448 resolution at 24 Frames Per Second (FPS). The camera recordings were triggered using ADS-B data by enforcing an 8 km fence around the setup. If an aircraft crosses the threshold, the camera automatically starts recording. The recording stops either at the end of a 250-sec timer or if the aircraft exits the 8 km fence. A new recording is activated if the recording times out before the aircraft exits. This process repeats till the aircraft leaves the geo-fence. To ensure seasonal diversity, data collection occurred in multiple phases from December 14, 2021, to February 23, 2023, from sunrise to sunset.

Post-Processing. As shown in Table 2, along with the raw camera output, TartanAviation offers various useful meta-data, including ADS-B and bounding box labels.

For ADS-B, each recorded video sequence has a PKL file that contains the GPS coordinates of all flying aircraft in the area at the specific time of recording. The ADS-B data also provides the unique flight ID and N-number (if provided by the pilot to the ADS-B) per aircraft. The unique flight ID and the N-number of the aircraft are used to parse the FAA N-number registry inquiry website. Furthermore, we obtain the ADS-B data and interpolate it to obtain ground truth information matching the frequency of the camera recording, i.e., 24 Hz.

For the bounding box labels, each recorded video sequence has a zip file containing the label files for individual video frames. The labels are text files, where each row provides the bounding box information for a flying aircraft in the video. In particular, the label files include the bounding box coordinates for the aircraft in the image, a unique track ID for each aircraft, the range of the aircraft, the manufacturer, type, and model.

We leverage two main modules to generate the bounding box labels: (a) Auto-labeling and (b) Manual label verification/generation. The Auto-labeling module is performed in two stages. The first stage uses our deep learning-based aircraft detection & tracking system, AirTrack⁹, to obtain initial bounding box hypotheses. Despite not having perfect recall, this stage makes labeling bounding boxes easier and reduces the amount of manual labor required downstream. The second stage projects the 3D ADS-B data into the image frame using the calibrated extrinsics to provide an initial estimate of the aircraft position in the image. This projection helps filter false positives from the detector and provides the initial airborne object location to the labeler. Once the auto-labeling procedure is finished, we manually verify the image labels frame-by-frame and make any necessary corrections using our custom labeling software.

The system used in the Auto-labeling module, AirTrack⁹, consists of modules for ego-aircraft motion estimation, detection and tracking, and secondary classification to filter out false positives. The inputs are two successive grayscale image frames $I_t, I_{t-1} \in \mathbb{R}^{H \times W}$ where $H \times W$ are the dimensions of the input frames. We utilize the full image resolution of 2048×2448 during inference to maximize the chances of detection at long ranges. The outputs are a list of tracked objects with bounding box coordinates, track ID, 2D Kalman filtered state, estimated range, angular rate, and time to closest point approach (tCPA). We then filter out false positive detections by projecting the ADS-B data into the image frame and ensuring that the projected ADS-B coordinates lie within a 100 pixel radius of the predicted image bounding box. Whenever we have false negatives from the model, we manually add in the bounding box label if the aircraft is visible in the image.

For the ADS-B data projection, we assume that the 4-camera setup is located at the origin of a north-east-down (NED) world reference frame where the positive z-axis points towards the earth's center and the positive x-axis points towards True North. Given all the relevant information, such as GPS coordinates and altitudes, we can compute the 3D position of aircraft using geodesics. The corresponding image frame location in homogeneous coordinates for camera i ($i = 1, \dots, 4$) is given by the following perspective projection: $\mathbf{p}_i = \mathbf{K}_i[\mathbf{R}_i|\mathbf{t}_i]\mathbf{P}$, where \mathbf{K}_i is a known 3×3 intrinsic matrix for camera i , $[\mathbf{R}_i|\mathbf{t}_i]$ constitutes the 3×4 extrinsic matrix of camera i , and $\mathbf{P} \in \mathbb{R}^3$ is the aircraft 3D position in the world reference frame. In particular, we use the PnP-RANSAC algorithm to estimate $[\mathbf{R}_i|\mathbf{t}_i]$. Using time synchronization, the 2D correspondences for the recorded 3D observations are manually labeled using custom labeling software. Since the setup is static, this is a one-time calibration step that gives us the projection matrix for camera i , with which we can project new 3D aircraft data.

Trajectory and Weather Data Acquisition. The trajectory and weather data are collected at both the KAGC and KBTP airports. Data collection follows similar procedures as our prior ADS-B dataset¹⁴. We use a Stratux ADS-B receiver capable of receiving position reports on both the 1090 MHz and 978 MHz frequencies. The receiver was installed within the terminal area of each airport. The recording operation ran from 1:00 AM to 11:00 PM local time, a period selected to encapsulate the full range of aviation operations during both peak and off-peak hours. Data collection at KBTP started on September 18, 2020 and concluded on October 27, 2022. Data collection at KAGC began on October 31, 2021 and ended on February 17, 2023. Data was collected in discrete phases, resulting in 381 days or 36 million raw position reports of data at KBTP and 280 days or 27 million raw position reports of data at KAGC. Weather reports for the corresponding airports are collected post hoc in the form of METeorological Aerodrome Reports (METAR) strings. We use the Iowa State METAR repository²¹ to compile reports for the duration of the ADS-B data. METAR reports are typically issued hourly, but the frequency can vary based on specific station protocols and prevailing weather conditions. Some stations release reports more frequently, such as three times per hour, especially when significant meteorological changes occur. Additionally, special reports, known as SPECIs, are generated in response to notable weather events or changes²². The total file size of the audio dataset is 1.9 GB compressed and 12 GB uncompressed.

Post-processing. The raw ADS-B data is first post-processed to remove corrupted and duplicate data points. The data is then filtered for altitude and distance from the airport. We nominally chose 6000 ft MSL and a 5 km radius around the airport. Once filtered, the data is transformed from a global to a local Cartesian coordinate frame in SI units with the origin at the end of the runway. The raw METAR strings are also processed to get wind velocity and direction relative to the runway in the local frame. Finally, we interpolate the trajectory data every second for all agents. TartanAviation provides the processed data along with the raw trajectory and weather data.

Speech Data. The speech data is collected at both the KAGC and KBTP airports. The setup uses a Bearcat SR30C radio receiver capable of receiving aviation radio frequencies. For towered KAGC, the radio is tuned to receive 121.1 MHz, the air traffic control tower frequency for KAGC operations. KAGC has an active tower 24 hours a day. For KBTP, the radio was tuned to 123.05 MHz, the Common Traffic Advisory Frequency for KBTP operations. The audio is recorded at a rate of 44.1k samples per second onboard the system. The record trigger mechanism for the speech setup is similar in construction to the vision setup and uses ADS-B data to trigger recordings. The trigger threshold was set to 10 km. The corresponding raw ADS-B data is also provided for the communication recordings spanning multiple years for both airports. The communication data from KBTP starts on September 6 2020 and ends on October 27 2022. After filtering, this data comprises 278 days of data, corresponding to 790 GB of disk space. The communications data for KAGC starts on October 31, 2021 and ends on February 17, 2023. The KAGC data comprises 392 days of data, corresponding to 1358 GB of disk space. In total, the dataset contains 670 days of raw communications recordings. The total file size of the audio dataset is 2.15 TBs uncompressed and 505.2 GB compressed.

Post-processing. We filtered the raw recordings using two conditions: audio clips must be longer than 1 second and have a maximum decibel level above -20 db. This removes erroneous recordings and clips that don't contain any spoken words. The KAGC audio has 33289 audio files comprising 2131.9 hours of audio, with 327.714 hours above -20 db. The KBTP audio has 8534 files and 1242.9 hours of audio, with 149.9 hours above -20 db. This gives a total of 41823 files, 3374.8 hours of audio, and 477.6 hours of audio above -20 db.

Data Records

The dataset²³ is stored on open access servers maintained by Carnegie Mellon University School of Computer Science and is available using the download scripts hosted on <https://github.com/castacks/TartanAviation>. [git\[https://doi.org/10.5281/zenodo.14699102\]](https://doi.org/10.5281/zenodo.14699102). We provide Python scripts to enable data download of each modalities, including support for sample and partial data. These scripts provide a platform-agnostic way to

Raw Data Fields			Processed Data Fields		
Field	Units	Description	Field	Units	Description
ID	#	ADS-B Aircraft ID	Frame	#	Relative Timestep
Time	HH:MM:SS.ss	Time of observation	ID	#	ADS-B Aircraft ID
Date	MM/DD/YYYY	Date of observation	x	km	Local X Cartesian Position
Altitude	Feet	Aircraft Altitude (Mean Sea Level)	y	km	Local Y Cartesian Position
Speed	Knots	Aircraft Speed	z	km	Local Z Cartesian Position
Heading	Degrees	Aircraft Heading	w_x	m/s	Component of wind along the dominant runway
Lat	Decimal Degrees	Latitude of the aircraft			
Lon	Decimal Degrees	Longitude of the aircraft	w_y	m/s	Component of wind across the dominant runway
Age	Seconds	Time since last observation			
Range	km	Distance from airport centre			
Bearing	Degrees	Bearing Angle with respect to North			
Tail		Aircraft Registration Number			
AltisGNSS	boolean	Indicator flag for Altitude Measurement			

Table 2. Variable description for the TartanAviation ADS-B trajectory data.




Fig. 3 Qualitative samples from the TartanAviation Image dataset showcasing the diversity of the collected images in different lighting conditions, seasons, cloud covers, and aircraft types.

download these datasets. Each of the modalities also have documentation that list any dependencies required to execute the scripts.

Additional information about the raw and processed data for all the modalities is available at <https://theair-lab.org/tartanaviation>. The following sections present details on the data formats and provide information on file organization.

Image Data. The image dataset is split across 550 independent sequences. We define a sequence as all of the data recorded during a single event where the camera recordings were started and stopped. The vision data folder contains multiple zipped files, each associated with a particular camera recording for that sequence. Each zipped sequence folder has multiple files, as presented in Table 1. Further information regarding the nomenclature and file contents is also shown in Table 1. In addition to the video files and labels, we also provide ADS-B data for each sequence.

Trajectory and Weather Data. TartanAviation provides both raw and processed data for each airport. Raw data is separated into individual folders for each day of collection. Each raw data folder has CSVs with fields detailed in Table 2. The processed files are available as comma-separated TXT files with fields described in 2.




Fig. 4 Log-normed trajectory histograms from ADS-B aircraft position reports at Allegheny County Airport.

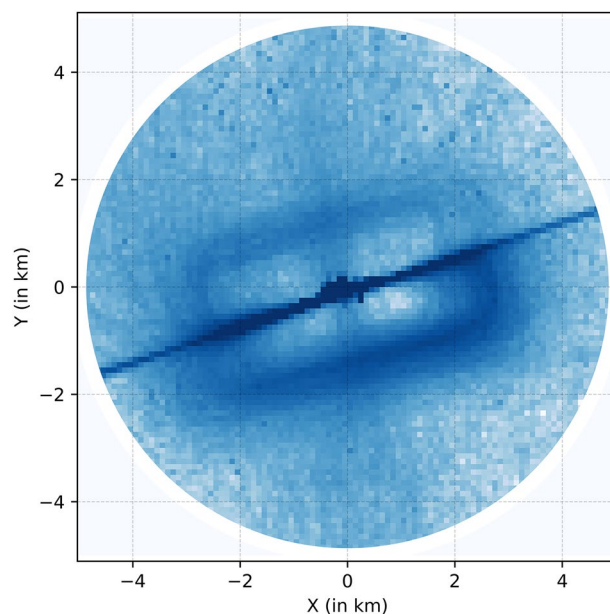


Fig. 5 Log-normed trajectory histograms from ADS-B aircraft position reports at Pittsburgh-Butler Regional Airport.

Speech Data. Both the raw and filtered audio files are included in the dataset. The filtered data is organized in a directory structure by location, year, month, and day. Each day is individually zipped, contains audio files in the WAV format, and has an accompanying text file that contains the audio clip's start, end, and total time.

Technical Validation

To ensure the reliability of our dataset, we employed tailored validation methods for each data modality: For ADS-B trajectory data, we conducted a comprehensive analysis by plotting and examining the entire processed dataset to identify any anomalies or inconsistencies. For speech data, we validated selected samples through detailed spectrogram analyses to ensure that vital speech frequencies were adequately captured. For image data, we performed both qualitative and quantitative assessments using selected samples and metadata from the entire dataset to showcase its diversity and quality. By implementing these modality-specific validation techniques, we ensured the overall integrity and robustness of our dataset.




Fig. 6 Quantitative diversity from the TartanAviation Image dataset showcasing the distribution of the collected images with respect to aircraft groups, cloud heights, and cloud cover.




Fig. 7 A portion of the spectrogram and waveform for audio file 6.wav from KBTP on November 2, 2020. The spectrogram shows the frequency of the audio signal versus time. In the time period of this figure, it is clear when the pilot is speaking, which can be seen as the higher intensity portion of the spectrogram and the spikes in amplitude in the waveform.

Image Data. Figure 3 shows example images from the dataset. The subset highlights the variation in lighting conditions, seasons, cloud covers, cloud heights, and aircraft types. The tightly labeled bounding boxes provided with each image are also shown. The scale of the bounding boxes with respect to the total image size highlights the challenge faced by object detection algorithms trying to reliably detect aircraft in a cluttered background. Figure 6 shows quantitative results for the entire dataset. To showcase the diversity in aircraft, we group the images based on single-engine land (SEL), multi-engine land (MEL), and Rotorcraft (rotor). This shows an almost equal distribution in the dataset with SEL as the major class with 56%. Further, we group the images based on the recorded cloud height. Cloud height has a direct impact on the available lighting. While 56.4% of the images have no clouds in them, 5.6% of images have very low cloud layers below 2000 feet above ground level (AGL). Grouping the images by cloud coverage, we observe that while 56.4% of data has no clouds, 12.2% of data has overcast skies. In addition to this, 6.4% of data has active precipitation while 12.4% of data has visibility less than 10 statute miles.

Trajectory Data. Figure 5 and Fig. 4 showcase trajectory histograms for the KAGC and KBTP Airports. The histograms represent the aircraft occupancy frequencies on a log scale around both airports. The airport is at the center in both images, with the geographic north pointing upwards. The effect of runway geometries is clearly visible. Figure 5 shows the 08/26 KBTP runway while Fig. 4 shows the crossing smaller 13/31 and larger 10/28 runways at KBTP. Also clearly visible are the different types of operations at both airports Figure 6. KBTP is an un-towered airport and home to a few flight schools. The left traffic patterns for both runways are clearly visible, highlighting the adherence to FAA guidelines when operating in an un-towered airfield. KAGC, on the other hand, is a towered airfield that hosts medical evacuation helicopters and business jet traffic in addition to flight schools. This leads to more straight-line arrivals and departures from the airport runways, as reflected in the figure.

Speech Data. Our filtering of the audio data ensures that empty files are discarded while any that could contain speech data are included. The threshold of -20 db was chosen to preserve any audio louder than radio silence. Figure 7 shows a portion of the spectrogram and waveform of one of the audio files from our dataset, specifically 6.wav from KBTP on November 2, 2020. It is apparent in the spectrogram where the pilot is speaking and where there is radio silence. Furthermore, some low-intensity noise during radio silence occurs at clear intervals and shows up as vertical lines with even spacing. During the speaking portion of the audio, there is a clear spike in intensity between all bands, with the highest being between 0-3 kHz. The human ear is highly sensitive in frequencies between 1-4 kHz, with the sensitivity dropping off steeply to the limit of around 20 kHz²⁴. Research in the interpretability of speech finds that the values below 4 kHz yield high accuracy in articulation²⁵. The spectrogram shows these vital frequencies for speech data are recorded sufficiently.

Usage Notes

In generating and post-processing TartanAviation, care was taken to organize data in formats commonly accepted by respective communities. This leads to hassle-free integration with existing dataloaders for specific tasks like vision-based small object detection and trajectory prediction. Pytorch Dataloaders for vision-based object detection are also provided in the supplementary codebase released with the manuscript. The data is available on high throughput servers without a paywall or any other access restrictions to enable widespread utilization. The provided download scripts ensure that the dataset can be downloaded in user-specified chunks to enable processing on local machines as well as large-scale servers. The modalities can be used independently or in conjunction. TartanAviation not only challenges the existing methods for established tasks like speech-to-intent prediction, object detection, and trajectory prediction but also opens up exciting avenues for leveraging multi-modal data in the context of aviation.

Code availability

The scripts to record, post-process and download each modality are publicly available at <https://github.com/castacks/TartanAviation.git>.

Received: 20 February 2024; Accepted: 7 March 2025;

Published online: 20 March 2025

References

- Administration, F. A. Faa aerospace forecast-fiscal years 2023-2043 (2023).
- Mueller, E. R., Kopardekar, P. H. & Goodrich, K. H. Enabling airspace integration for high-density on-demand mobility operations. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, 3086 (2017).
- Association, I. A. T. Air passenger market analysis - december 2022. Tech. Rep., IATA (2022).
- Lin, Y. *et al.* Towards recognition for radio-echo speech in air traffic control: Dataset and a contrastive learning approach. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2023).
- Gomez, J. P. Z. *et al.* ATCO2 corpus: A Large- Scale Dataset for Research on Automatic Speech Recognition and Natural Language Understanding of Air Traffic Control Communications. *Journal of Data-centric Machine Learning Research* (2024).
- Smyers, E., Katz, S., Corso, A. & Kochenderfer, M. J. Avoidds: Aircraft vision-based intruder detection dataset and simulator. *Advances in Neural Information Processing Systems*, **36**, 8076-8091 (2023).
- AIcrowd. Airborne object tracking challenge (2022).
- Weibel, R. & Hansman, R. J. Safety considerations for operation of different classes of uavs in the nas. In *AIAA 4th Aviation Technology, Integration and Operations (ATIO) Forum*, 6244 (2004).
- Ghosh, S., Patrikar, J., Moon, B., Hamidi, M. M. & Scherer, S. Airtrack: Onboard deep learning framework for long-range aircraft detection and tracking. <https://doi.org/10.48550/ARXIV.2209.12849> (2022).
- Farhadmanesh, M., Rashidi, A. & Marković, N. General aviation aircraft identification at non-towered airports using a two-step computer vision-based approach. *IEEE Access* **10**, 48778-48791 (2022).
- Martin, J., Riseley, J. & Ford, J. J. A dataset of stationary, fixed-wing aircraft on a collision course for vision-based sense and avoid. In *2022 International Conference on Unmanned Aircraft Systems (ICUAS)*, 144-149 (IEEE, 2022).
- Ali, B. *S.A safety assessment framework for automatic dependent surveillance broadcast (ads-b) and its potential impact on aviation safety*. Ph.D. thesis, Centre for Transport studies, Department of Civil and Environmental - (2013).
- Xue, D., Hsu, L.-T., Wu, C.-L., Lee, C.-H. & Ng, K. K. *Advanced Engineering Informatics* **50**, 101402 (2021).
- Patrikar, J., Moon, B., Oh, J. & Scherer, S. 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, 2022).
- Xue, D., Yang, J., Liu, Z. & Yu, S. Examining the economic costs of the 2003 halloween storm effects on the north hemisphere aviation using flight data in 2019. *Space Weather* **21**, e2022SW003381 (2023).
- Godfrey, J. J. Air traffic control complete. *Linguistic Data Consortium, Philadelphia, USA* <http://www ldc.upenn.edu/Catalog/CatalogEntry.jsp> (1994).
- Guo, D. *et al.* M2ats: A real-world multimodal air traffic situation benchmark dataset and beyond. In *Proceedings of the 31st ACM International Conference on Multimedia*, 213-221 (2023).
- Srinivasamurthy, A. *et al.* Semi-supervised learning with semantic knowledge extraction for improved speech recognition in air traffic control. Tech. Rep. (2017).
- U.S., Congress 18 U.S. Code § 2511 - Interception and disclosure of wire, oral, or electronic communications prohibited, <https://www.law.cornell.edu/uscode/text/18/2511> (2025).
- Federal Aviation Administration ADS-B Privacy, https://www.faa.gov/air_traffic/technology/equipadsb/privacy (2025).
- Herzmann, D., Arritt, R. & Todey, D. Iowa environmental mesonet. Available at mesonet.agron.iastate.edu/request/coop/fe.phtml (verified 27 Sept. 2005). *Iowa State Univ., Dep. of Agron., Ames, IA* (2004).
- National Weather Service National Weather Service Observing Handbook No. 8: Aviation Weather Observations for Supplementary Aviation Weather Reporting Stations (SAWRS), <https://www.weather.gov/media/surface/WSOH8.pdf> (1996).
- Patrikar, J. *et al.* castacks/tartanaviation: 1.0, <https://doi.org/10.5281/ZENODO.14699102> (2025).
- Sivian, L. J. & White, S. D. On minimum audible sound fields, <https://doi.org/10.1121/1.1915608> (1933).
- French, N. R. & Steinberg, J. C. Factors governing the intelligibility of speech sounds, <https://doi.org/10.1121/1.1916407> (1947).

Acknowledgements

This work is supported by the Mitsubishi Heavy Industries (MHI) project #A025279. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745016. We would like to thank the Allegheny County Airport Authority and the Office of Manager of Pittsburgh-Butler Regional Airport for their support.

Author contributions

The research project was conceived by S.S. The experiment was designed by J.P. and S.G. The setup was maintained by J.P., M.H., B.M., and I.H. The data was processed and results were analysed by J.P., N.K., B.M., M.H., J.D., and A.C. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to J.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025