

Operational Contextualization and Validation of AI-Based 4D Trajectory Predictions

Lucía Bragado ^{*}, Carlos Domínguez-Monferrer , and Álvaro Quintanar 

Indra, Madrid, Spain

*Corresponding author: lbragado@indra.es

Abstract

AI-based 4D trajectory prediction models generate spatiotemporal flight paths as sequences of latitude, longitude, altitude, and timestamp samples. Although these raw outputs provide detailed trajectory information, they are often difficult to interpret operationally because they do not explicitly reveal route intent, known waypoints, airways, traffic-flow structures, or external constraints affecting the predicted path. As a result, prediction accuracy alone may not provide sufficient evidence for operational understanding, plausibility assessment, or user trust. This paper investigates how AI-predicted 4D trajectories can be operationally contextualized and validated to support TBO.

Building on a previously published LSTM and Transformer-based 4D trajectory predictor, formulated as a deterministic sequence-to-sequence regression task, we present a proof-of-concept visual analytics framework that operates strictly on the model output. The framework enriches AI-generated trajectories with aviation-context layers through two complementary functions: contextualization, through route-context matching, waypoint proximity inference, and similarity assessment against known route structures; and validation, through spatial-temporal checks against active SIGMETs and NOTAMs that flag intersections only when constraint activity windows coincide with predicted aircraft passage.

A prototype is demonstrated in various pairs of ADEP-ADES, where the predicted 4D trajectory is contextualized against alternatives to FAA Preferred Routes, Coded Departure Routes, and operational constraint layers. The results suggest that operational contextualization transforms raw AI trajectory output into explainable trajectory intelligence, providing a human-centred and interpretable bridge between data-driven prediction and TBO. By presenting historical and contextual evidence without modifying the prediction or implementing operational actions, the framework supports human assessment while preserving full operator authority. The proposed approach is transferable to other airspaces where equivalent operational data sources are available.

Keywords: 4D trajectory prediction; Trajectory-Based Operations; Deep Learning; Clustering; Traffic flow

Abbreviations: TBO: Trajectory-Based Operations, FAA: Federal Aviation Administration, CDR: Coded Departure Route, SIGMET: Significant Meteorological Information, NOTAM: Notice to Airmen

1. Introduction

Four-dimensional (4D) trajectory-prediction models estimate future aircraft positions as sequences of latitude, longitude, altitude, and time values. Although deep-learning models can generate accurate spatiotemporal trajectories, their numerical outputs do not directly indicate whether a prediction follows a known route, resembles recurrent traffic patterns, or potentially interacts with active weather phenomena or airspace restrictions.

Trajectory prediction also presents an intrinsic data-representativeness challenge. Large historical datasets do not necessarily provide sufficient operational diversity. For many origin-destination pairs, most flights follow a small number of dominant route structures, causing machine-learning

models to favour recurrent patterns while potentially underrepresenting less frequent but operationally relevant alternatives. Consequently, geometric prediction accuracy alone provides limited evidence of a trajectory’s operational plausibility and relevance.

In previous work, we developed the HFP trajectory-prediction capability using Long Short-Term Memory (LSTM) and Transformer architectures [1]. This paper extends HFP with a *model-agnostic post-prediction layer* that contextualizes the generated trajectory. The framework combines four sources of operational context (see Figure 1): FAA Preferred Instrument Flight Rules (IFR) Routes and Coded Departure Routes (CDRs), historical traffic patterns, Significant Meteorological Information (SIGMET) advisories and weather-radar reflectivity, and Temporary Flight Restrictions (TFRs) published through Notices to Air Missions (NOTAMs).




Figure 1. Overview of the HFP architecture and its main layers.

Rather than relying exclusively on geometric prediction accuracy, the framework supplies additional evidence that helps users understand how a predicted trajectory relates to published route structures, recurrent traffic patterns, weather phenomena, and active airspace restrictions.

Although the prototype uses data from the United States National Airspace System, the proposed approach is transferable to other airspaces where equivalent operational data sources are available.

2. Methods

2.1 Framework and Data

The framework is implemented as a client–server application operating after model inference. A predicted trajectory is represented as

$$P = \{(\text{lat}_i, \text{lon}_i, \text{alt}_i, t_i)\}_{i=1}^N. \quad (1)$$

The original sequence remains unchanged, while additional analytical layers are incorporated to provide operational context and plausibility evidence. Published routes (IFR and CDRs) and historical trajectories obtained through SWIM are used for contextualization purposes, allowing the prediction to be interpreted against known traffic structures and operational practices. Meteorological and airspace restrictions are used for validation purposes, identifying potential interactions between the prediction and active operational constraints.

2.2 Trajectory Generation

The HFP LSTM and Transformer models generate future sequences of latitude, longitude, altitude, and timestamps from flight and observed-state information. The trained models are exported to

ONNX format and integrated into the HFP server to enable real-time inference within the contextualization framework.

Since the focus of this work is the post-prediction contextualization and validation layer, further details regarding model architecture, training procedures, dataset partitioning, and trajectory-prediction performance are available in the previous trajectory-prediction study.

2.3 Route and Historical-Pattern Contextualization

For each ADEP-ADES pair, applicable published routes are retrieved and converted from symbolic descriptions into geolocated waypoint sequences. Airways are expanded using NASR data, while ambiguous elements are resolved according to route order and geographical consistency.

Historical trajectories are grouped using HDBSCAN and represented by cluster centroids. The similarity between a predicted trajectory and each centroid is quantified through the area enclosed between both trajectories. The centroid yielding the minimum area is selected as the closest historical traffic pattern and reported as contextual information associated with the prediction.

2.4 Constraint Interaction Assessment

SIGMET advisories and TFR NOTAMs are represented as georeferenced polygons associated with validity intervals. Potential interactions are identified by evaluating the spatial intersection between the predicted trajectory and each constraint and comparing the estimated trajectory passage time against the corresponding validity period.

When altitude limits are available, vertical compatibility is also assessed. Constraints with incomplete temporal or vertical information are displayed to the user but are marked as only partially assessed.

NOAA radar reflectivity is incorporated as a complementary observational layer and is not currently used within the automated interaction-assessment process.

Meteorological and airspace information is used exclusively to identify operational factors that may affect the plausibility of the predicted trajectory. The framework neither modifies the predicted trajectory nor generates alternative routing solutions.

2.5 Experimental Protocol

The prototype is evaluated on a set of representative ADEP-ADES pairs. Following the modeling approach adopted in previous work, an independent trajectory prediction model is trained for each airport pair, allowing the predictor to capture route-specific operational patterns and trajectory dynamics. For each evaluated flight, the observed trajectory, the nearest published route, and the closest historical centroid are used as operational references. The analysis considers lateral, vertical, and temporal deviations, together with potential interactions with active operational constraints. Visual examples complement the quantitative comparison.

3. Results

Figure 2 illustrates the framework applied to a DAL A321 flight operating between Miami (KMIA) and Atlanta (KATL) on 31 July 2026 at 22:11 UTC. The HFP model generated an 81-point 4D trajectory prediction, subsequently enriched with operational context layers.

For this case, the framework retrieved five historical traffic patterns and incorporated FAA published-route references, four active SIGMET advisories, TFR NOTAMs, and NOAA radar reflectivity. The resulting view supports comparison of the prediction with recurrent traffic structures and published

routes, while enabling the spatial, temporal, and, when available, vertical assessment of potential interactions with active constraints.

The example demonstrates the main contribution of the framework: transforming a sequence of predicted spatiotemporal points into an operationally interpretable representation. The framework connects the prediction with route-level information, historical behavior, and constraint-interaction evidence that cannot be obtained from the model output alone.




Figure 2. Visual analytics prototype applied to the Miami–Atlanta corridor.

4. Conclusions

This paper presented a framework for the operational contextualization of AI-based 4D trajectory predictions. By combining published routes, historical traffic patterns, meteorological information, and airspace restrictions, the framework complements raw trajectory predictions with operational evidence that helps users interpret their plausibility and relevance.

The framework follows a human-centred approach in which the AI-generated trajectory and its associated contextual evidence support human assessment without selecting or implementing operational actions. Under the EASA AI classification, the current prototype is consistent with Level 1B human assistance principles, with the end user retaining full authority and responsibility for any resulting operational decision.

The proposed post-prediction explainability approach is also aligned with the relevance principles described in EASA anticipated MOC EXP-04-2. In particular, historical traffic patterns provide a historical perspective, while published routes, meteorological conditions, airspace restrictions, and their corresponding data sources provide contextual evidence that can help the user assess the operational relevance of the prediction. This constitutes output-level operational explainability rather than an explanation of the internal reasoning of the underlying LSTM or Transformer model.

Author contributions

- Lucía Bragado: Conceptualization, Methodology, Software, Data Curation, Resources, Writing – Original Draft, Writing – Review and Editing.
- Carlos Domínguez-Monferrer: Validation, Formal Analysis, Investigation, Visualization, Resources, Writing – Review and Editing.
- Álvaro Quintanar: Supervision, Resources, Writing – Review and Editing.

Open data statement

This work combines data obtained from publicly available aviation and meteorological sources, including FAA NASR navigation data, FAA Preferred Routes, Coded Departure Routes (CDRs), AWC SIGMET advisories, FAA TFR notices, and NOAA weather products. Historical trajectory information used to derive traffic patterns was obtained from SWIM-based operational data sources and is therefore subject to the corresponding access and distribution restrictions. Further information regarding the data sources used in this study can be provided by the authors upon reasonable request.

Reproducibility statement

The methodological workflow and data sources used in this study are described in sufficient detail to allow independent implementation. The framework relies on publicly available aviation and meteorological information sources, including FAA Preferred Routes, FAA Coded Departure Routes, FAA Temporary Flight Restriction (TFR) notices, Aviation Weather Center (AWC) SIGMET products, and NOAA weather data. The software prototype and HFP implementation used in this work are proprietary and therefore cannot be publicly released. However, the public data sources and acquisition endpoints employed in the study are documented and can be independently accessed by other researchers.

References

- [1] Álvaro Quintanar, Lucía Bragado, and Hamza Triki. “Deep Learning Framework for 4D Trajectory Prediction: From Airport Pairs to Full Flight Paths”. In: *2026 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. 2026, pp. 1–14. DOI: [10.1109/ICNS69853.2026.11570276](https://doi.org/10.1109/ICNS69853.2026.11570276).
- [2] Federal Aviation Administration. *Preferred IFR Routes Database*. 2026. URL: https://www.fly.faa.gov/rmt/data_file/prefroutes_db.csv.
- [3] Federal Aviation Administration. *Coded Departure Routes Database*. 2026. URL: https://www.fly.faa.gov/rmt/data_file/codedswap_db.csv.
- [4] Aviation Weather Center. *Aviation Weather Data API*. 2026. URL: <https://aviationweather.gov/api/data>.
- [5] National Oceanic and Atmospheric Administration. *Radar Base Reflectivity*. 2026. URL: https://mapservices.weather.noaa.gov/eventdriven/rest/services/radar/radar_base_reflectivity/MapServer.
- [6] Federal Aviation Administration. *Temporary Flight Restrictions (TFR)*. 2026. URL: <https://tfr.faa.gov>.