Proceedings of the Thirty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-26)
Demonstrations Track

ORBIT: Optimal Recommendation Framework for Boarding with Interpretable Timelines

Joonseong Kang[1,2,*], **Seung Ha Hwang**[1,3,*],
Jaehun Bang[1,4,*], **Jiyoung Ko**[1,5,*], **Subeen Park**[1,2,†] and **Jeffrey Gennari**[1,†]
[1]Carnegie Mellon University
[2]Yonsei University
[3]Kyung Hee University
[4]UNIST
[5]Jeju National University
{joonseok, jaehunb, seunghah, jko3, subeenp, jgennari}@andrew.cmu.edu

Abstract

Determining when to leave for the airport is a complex problem shaped by flight delay risk, traffic, weather, airport congestion, and passenger-specific constraints. Existing services rely on isolated delay estimates or simple travel time calculations, failing to capture real-time context. We propose ORBIT, a decision-support system that generates personalized leave-by recommendations by integrating predictive models with real-time operational signals. The system combines user input normalization, real-time data acquisition, Transformer-based delay prediction, and LLM-based reasoning to jointly account for statistical forecasts and dynamic factors like congestion, previous-leg propagation, weather, traffic, and airport processing times. Instead of a standalone delay estimate, ORBIT produces an actionable and interpretable departure plan. We implement ORBIT as an interactive system and demonstrate its applicability. A video demo is available at https://youtu.be/fZX42jceIM4.

1 Introduction

Existing airport departure planning services typically rely on single-source static estimation or simple travel time calculation, and therefore do not fully capture real-time operational context or personalized mobility constraints [Rebollo and Balakrishnan, 2014]. In practice, the optimal leave-by time is jointly determined by flight delay risk, traffic conditions, weather factors, airport congestion, and passenger-specific characteristics [Small *et al.*, 2005]. Considering any one of these components in isolation is insufficient for reliable decision support in real-world travel scenarios.

To address these limitations, we propose ORBIT, a decision-support system that generates personalized leave-by recommendations by jointly considering flight delay factors and integrated mobility context, including traffic, weather, and airport operations [Pyrgiotis *et al.*, 2013; Sternberg *et al.*, 2017]. ORBIT is structured as a multimodal pipeline combining user input processing, real-time data acquisition and normalization, gemini-based flight delay regression [Team *et al.*, 2024], and large language model (LLM)-based reasoning integration. Instead of providing a standalone delay estimate, the system delivers an actionable departure plan that accounts for both predictive uncertainty and operational constraints.

The contributions of this work are threefold. First, we integrate Transformer-based delay prediction [Vaswani *et al.*, 2017] with real-time signals and LLM explanations. Second, we personalize ground access and airport processing using user location, transportation mode, and baggage profile [Koster *et al.*, 2011]. Third, we provide an explainable interface that decomposes leave-by recommendations into temporal components [Doshi-Velez and Kim, 2017].

2 System Design

The system follows a four-stage pipeline: USER INPUT → DATA COLLECTION & PREPROCESSING → DELAY PREDICTION → LLM INTEGRATION. The objective of the system is to compute a recommended leave-by time by jointly accounting for the likelihood of flight departure delay and the time required for ground travel and airport-side procedures. We designated John F. Kennedy International Airport (JFK) in New York as the primary target airport, as it is a high-traffic hub with diverse operational conditions.

2.1 User Input Collection

The system collects the following user inputs to produce a personalized recommendation: origin address, flight number, airline code, departure and arrival airports (International Air Transport Association, IATA codes), scheduled departure time, transportation mode (Driving/Transit/Walk/Bicycle), checked baggage status, and terminal/gate information.

User inputs are supported through three channels: Text-based interface entry, Image-based ticket optical character recognition (OCR), and Speech-based input. Ticket images are processed using Gemini-based OCR [Team *et al.*, 2024], and voice input is transcribed using faster-whisper [Radford *et al.*, 2023]. The extracted information is unified into a structured representation for downstream modules.

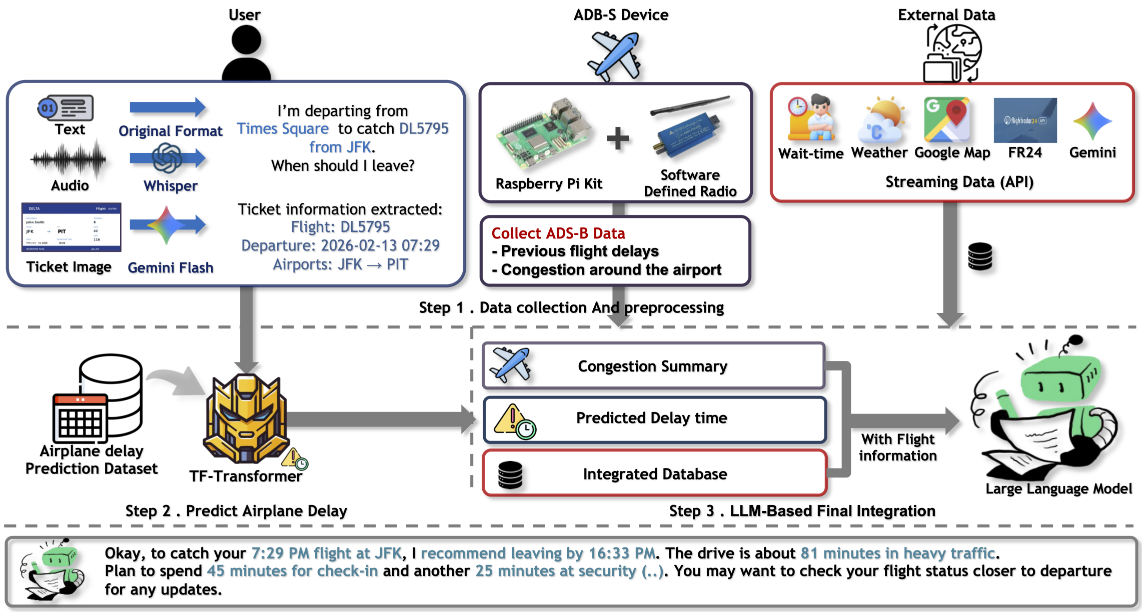

Figure 1: **ORBIT architecture pipeline.** Users input departure and flight information via text, voice, or boarding pass images. The ADS-B device collects real-time congestion and delay data via Raspberry Pi and SDR. External APIs provide wait times, weather, and flight status. A TF-Transformer model predicts delays, and an LLM generates leave-by time recommendations with an interpretable rationale.

2.2 Data Collection and Preprocessing

This study uses both offline training data and online operational data. For offline training, we construct delay-prediction features from the Kaggle Airplane Delay Dataset.[1] The core features include airline identity, origin and destination airports, and temporal variables. During online operation, the system collects the following external data sources:

1. **Regional congestion near the departure airport ($T_{\text{congestion}}$).** Automatic dependent surveillance-broadcast (ADS-B) data around JFK are collected for multiple days to compute mean and standard deviation statistics [Yang *et al.*, 2023; Schäfer *et al.*, 2014]. At runtime, current traffic levels are compared against these statistics, and additional congestion time is applied when traffic exceeds the baseline.
2. **Previous-leg delay propagation using ADS-B and FR24 (T_{pflight}).** Using ADS-B data, the system estimates remaining flight time based on haversine distance and an assumed 450 kt cruise speed [Yang *et al.*, 2023; Andreou *et al.*, 2023]. The discrepancy between expected and observed elapsed time is modeled as propagated [Fleurquin *et al.*, 2013].
3. **Weather conditions (Google Weather) (T_{weather}).** The Google Weather API is used to assign a weather risk level, and the system adds an additional delay buffer of 0, 15, or 30 minutes accordingly.
4. **Travel time and route (Google Routes) (T_{traffic}).** The Google Routes API is used to estimate travel time from the user's origin to the airport, conditioned on the selected transportation mode.
5. **TSA wait time and gate walking time (T_{airport}).** TSA (Transportation Security Administration) wait time, walking time, and security-related timing information are retrieved from the JFK airport website.[2]

Because ORBIT depends on multiple external sources, it implements caching and fallback strategies to ensure operational robustness. When an API call fails, the system serves stale cache entries when available, or resorts to historical statistics and conservative defaults (e.g., travel time, TSA wait time, weather buffer, and delay estimates) so that the pipeline continues to produce an output.

We define the integrated external time factor as:

$$T_{\text{external}} = (T_{\text{weather}} + T_{\text{traffic}} + T_{\text{airport}}) - (T_{\text{congestion}} + T_{\text{pflight}}) \quad (1)$$

2.3 Transformer-based Delay Prediction

This study employs the TF-Transformer [Xiong *et al.*, 2020] to perform regression-based prediction of flight delays (in minutes). Each feature is tokenized and passed through self-attention blocks. The representation of the `[CLS]` token is then fed into a regression head to estimate the predicted delay time. During deployment, the delay estimation logic is defined as follows. The TF-Transformer prediction output is used to compute T_{tpredict} [Gorishniy *et al.*, 2021].

2.4 LLM-based Final Integration

In the Streamlit-based implementation, the airport process timeline is computed by backward calculation from the ac-

[1] https://www.kaggle.com/datasets/shubhamsingh42/flight-delay-dataset-2018-2024

[2] https://www.jfkairport.com/

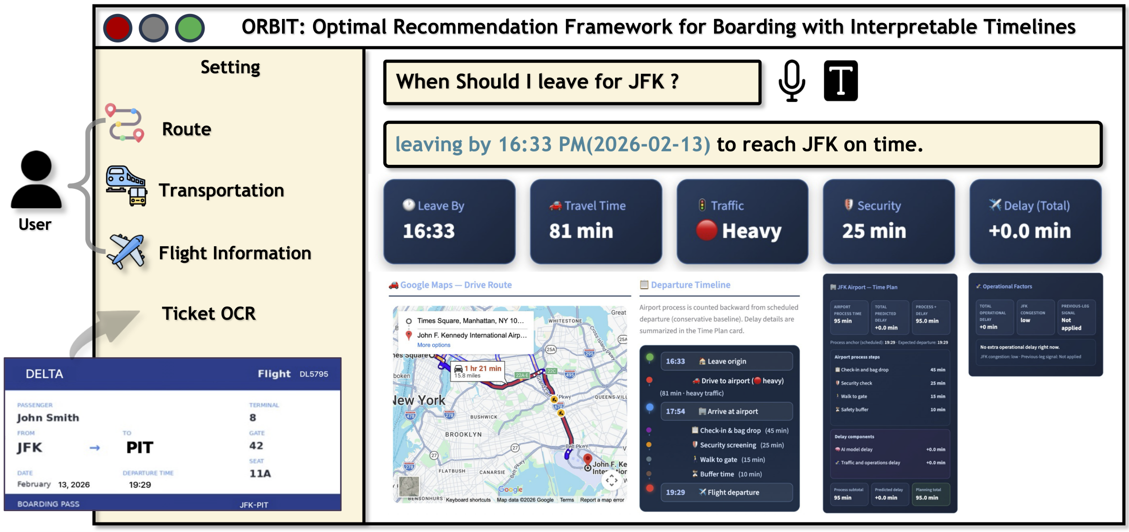

Figure 2: **ORBIT demonstration UI.** Users input their origin, transportation mode, and flight information, and the system presents a recommended leave-by time. Simultaneously, it decomposes traffic congestion, estimated travel time, and airport process times by component to provide an interpretable rationale. The lower section visualizes a map-based route and timeline, with expandable sections for detailed calculations and input summary. Further details are available in the demonstration video.

tual departure time.

The predicted departure time T_{leave} is defined as:

$$T_{\text{leave}} = T_{\text{actual}} - T_{\text{fixed}} - T_{\text{external}} + T_{\text{tpredict}} \tag{2}$$

where

- T_{actual}: Actual scheduled departure time.
- T_{fixed}: Predefined airport arrival requirement.
- T_{external}: Integrated external time factor.
- T_{tpredict}: Predicted delay from the TF-Transformer.

Here, T_{fixed} denotes the target airport arrival time (2 hours before departure). It includes a user-adjustable buffer and optional baggage drop time. The extracted trip constraints and computed time components are formatted into a structured prompt, which is provided to the LLM to generate the final leave-by recommendation and explanation [Yao *et al.*, 2022].

3 Implementation

The core output of the framework is the recommended leave-by time, T_{leave}. Rather than providing a single timestamp, the system presents a visual, itemized decomposition of the major time components to clarify the rationale and dominant factors behind the recommendation. Our code is available at https://github.com/T-Rex-iitp/ORBIT.

3.1 Dashboard

The Streamlit-based dashboard integrates multimodal inputs with explainable outputs. It supports three input modalities: (1) manual text entry of flight information, (2) boarding-pass or ticket image upload with OCR-based automatic extraction, and (3) in-browser voice input with real-time transcription (e.g., "When should I leave to JFK?"). Upon receiving input, the system extracts key entities such as the flight identifier, airport code, and relevant times, then invokes the recommendation engine to generate a structured departure brief.

The dashboard outputs consist of: (i) a visual departure timeline that decomposes the total preparation time into components, (ii) an interactive map displaying the origin-to-airport route and, when available, real-time contextual indicators, and (iii) an airport intelligence panel summarizing terminal and gate context, estimated security screening wait time, and current weather conditions. The system also generates a natural-language recommendation that translates the time-budget components into actionable guidance.

3.2 Demo Scenario

In a typical demo run for a passenger departing from JFK, flight information is provided via ticket image upload, manual entry, or voice. The system then (1) estimates the expected delay using an TF-Transformer, yielding a predicted delay component T_{tpredict}, (2) incorporates operational context such as congestion $T_{\text{congestion}}$ and previous-leg effects T_{pflight}, (3) computes origin-to-airport travel time T_{traffic}, (4) applies airport process times T_{airport} including security screening and terminal or gate transfer, and (5) reflects current weather conditions T_{weather}. Finally, the user is presented with the recommended leave-by time, a natural-language explanation, and an itemized breakdown of the contributing components.

4 Conclusion

This study proposed ORBIT, a multimodal decision support system that addresses flight departure uncertainties and individual mobility constraints. ORBIT integrates Transformer-based delay prediction with real-time operational signals to compute actionable leave-by times. By leveraging LLMs to synthesize heterogeneous temporal factors, the system ensures transparent and interpretable recommendations.

Ethical Statement

The proposed system processes user location data locally for inference and does not store or retain personally identifiable information. When external services (e.g., Google Maps or Weather APIs) are invoked, only non-identifying or coarse-grained parameters are transmitted whenever feasible, and no raw user inputs are persisted beyond the active session. The ADS-B data utilized in this work is publicly available aviation data and is used in compliance with applicable regulations.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (RS-2022-00143911, AI Excellence Global Innovative Leader Education Program)

Contribution Statement

Joonseong Kang, Seung Ha Hwang, Jaehun Bang, and Jiyoung Ko contributed equally to this work. Subeen Park and Jeffrey Gennari are corresponding authors.

References

[Andreou *et al.*, 2023] Andreas Andreou, Constandinos X Mavromoustakis, Jordi Mongay Batalla, Evangelos K Markakis, George Mastorakis, and Shahid Mumtaz. Uav trajectory optimisation in smart cities using modified a* algorithm combined with haversine and vincenty formulas. *IEEE Transactions on Vehicular Technology*, 72(8):9757–9769, 2023.

[Doshi-Velez and Kim, 2017] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.

[Fleurquin *et al.*, 2013] Pablo Fleurquin, José J Ramasco, and Victor M Eguiluz. Systemic delay propagation in the us airport network. *Scientific reports*, 3(1):1159, 2013.

[Gorishniy *et al.*, 2021] Yury Gorishniy, Ivan Rubachev, Valentin Khrulkov, and Artem Babenko. Revisiting deep learning models for tabular data. *Advances in neural information processing systems*, 34:18932–18943, 2021.

[Koster *et al.*, 2011] Paul Koster, Eric Kroes, and Erik Verhoef. Travel time variability and airport accessibility. *Transportation Research Part B: Methodological*, 45(10):1545–1559, 2011.

[Pyrgiotis *et al.*, 2013] Nikolas Pyrgiotis, Kerry M Malone, and Amedeo Odoni. Modelling delay propagation within an airport network. *Transportation Research Part C: Emerging Technologies*, 27:60–75, 2013.

[Radford *et al.*, 2023] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR, 2023.

[Rebollo and Balakrishnan, 2014] Juan Jose Rebollo and Hamsa Balakrishnan. Characterization and prediction of air traffic delays. *Transportation research part C: Emerging technologies*, 44:231–241, 2014.

[Schäfer *et al.*, 2014] Matthias Schäfer, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. Bringing up opensky: A large-scale ads-b sensor network for research. In *IPSN-14 proceedings of the 13th international symposium on information processing in sensor networks*, pages 83–94. IEEE, 2014.

[Small *et al.*, 2005] Kenneth A Small, Clifford Winston, and Jia Yan. Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica*, 73(4):1367–1382, 2005.

[Sternberg *et al.*, 2017] Alice Sternberg, Jorge Soares, Diego Carvalho, and Eduardo Ogasawara. A review on flight delay prediction. *arXiv preprint arXiv:1703.06118*, 2017.

[Team *et al.*, 2024] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[Xiong *et al.*, 2020] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International conference on machine learning*, pages 10524–10533. PMLR, 2020.

[Yang *et al.*, 2023] Zhanji Yang, Xiaolei Kang, Yuanhao Gong, and Jiansheng Wang. Aircraft trajectory prediction and aviation safety in ads-b failure conditions based on neural network. *Scientific reports*, 13(1):19677, 2023.

[Yao *et al.*, 2022] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*, 2022.