Flight Delay Analytics & Prediction Final Poster

Introduction/Motivation

What's the problem?
Flight delays cause missed connections, added costs, and cascading disruptions across the national air network. Small delays on a single aircraft can propagate to later flights through scheduling and weather impacts.

Why does this matter?
Millions of U.S. flights operate yearly—minor disruptions translate into system-wide effects. Accurate, real-time delay prediction supports airlines (operational scheduling, tail rotation decisions) and helps passengers select lower-risk itineraries.

Method

What are our approaches?
We develop a production-ready prediction pipeline combining:
•Lightweight Logistic Regression (final model)
•Tuned XGBoost (comparison benchmark)
•Full-feature Logistic Regression (baseline)
•Real-time inference via Tableau + TabPy
•SHAP-style natural-language feature explanations

Why will this work?
•Delays propagate through aircraft rotations → previous-leg delay features
•Weather drives seasonal spikes → rolling weather anomalies
•Route-level reliability captures structural congestion
•Lightweight model enables live inference + interpretability

Key innovations
•Propagation-aware tail-history features
•Real-time inference pipeline using NOAA weather feeds
•3-stage dashboard linking historical context → live predictions
•Compact 15-feature model optimized for real-time deployment

Data & Feature Engineering

How did we get it?
•BTS On-Time Performance (schedule, delay, tail-number histories)
•NOAA airport weather (temperature, wind, precipitation, snow)
•Joined via airport + timestamp alignment
•Stored in Parquet + aggregated via Python / Spark
Characteristics
•Tens of millions of flight records
•Hourly weather + rolling multi-hour summaries
•Multiple regions / seasons → generalization across hubs
•Rolling rate features computed only from train windows (no leakage)
•Key feature groups
•Weather: rolling 3/6/12/24-hr stats, anomalies
•Temporal: hour, month (cyclic encoding)
•Propagation: previous-leg delay signals
•Congestion: rolling traffic counts
•Historical reliability: airport + route delay rates

Evaluation - Model Results

How we evaluated
•Season-aware rolling time split (train past 3 weeks → test next week)
•Prevents temporal leakage, simulates real forecasting conditions
•Metrics: ROC-AUC, precision, recall, F1, calibration

Results (high-level)
•XGBoost performs best but higher latency + less interpretable
•Compact Logistic Regression nearly matches baseline while supporting live inference
•Weather anomaly + previous-leg features are most predictive

Interpretability
•SHAP explanations highlight: previous-leg delay, congestion, temp anomalies, airport reliability.

Model	ROC-AUC	Recall (delay)	Notes
LR(full)	0.789	0.663	Baseline
XGBoost	0.813	0.639	Best Perfromance
LR(15-feature)	0.783	0.619	Deployed (real-time)

Evaluation - Dashboard & Visualization Layer

Origin Airport Exploratory

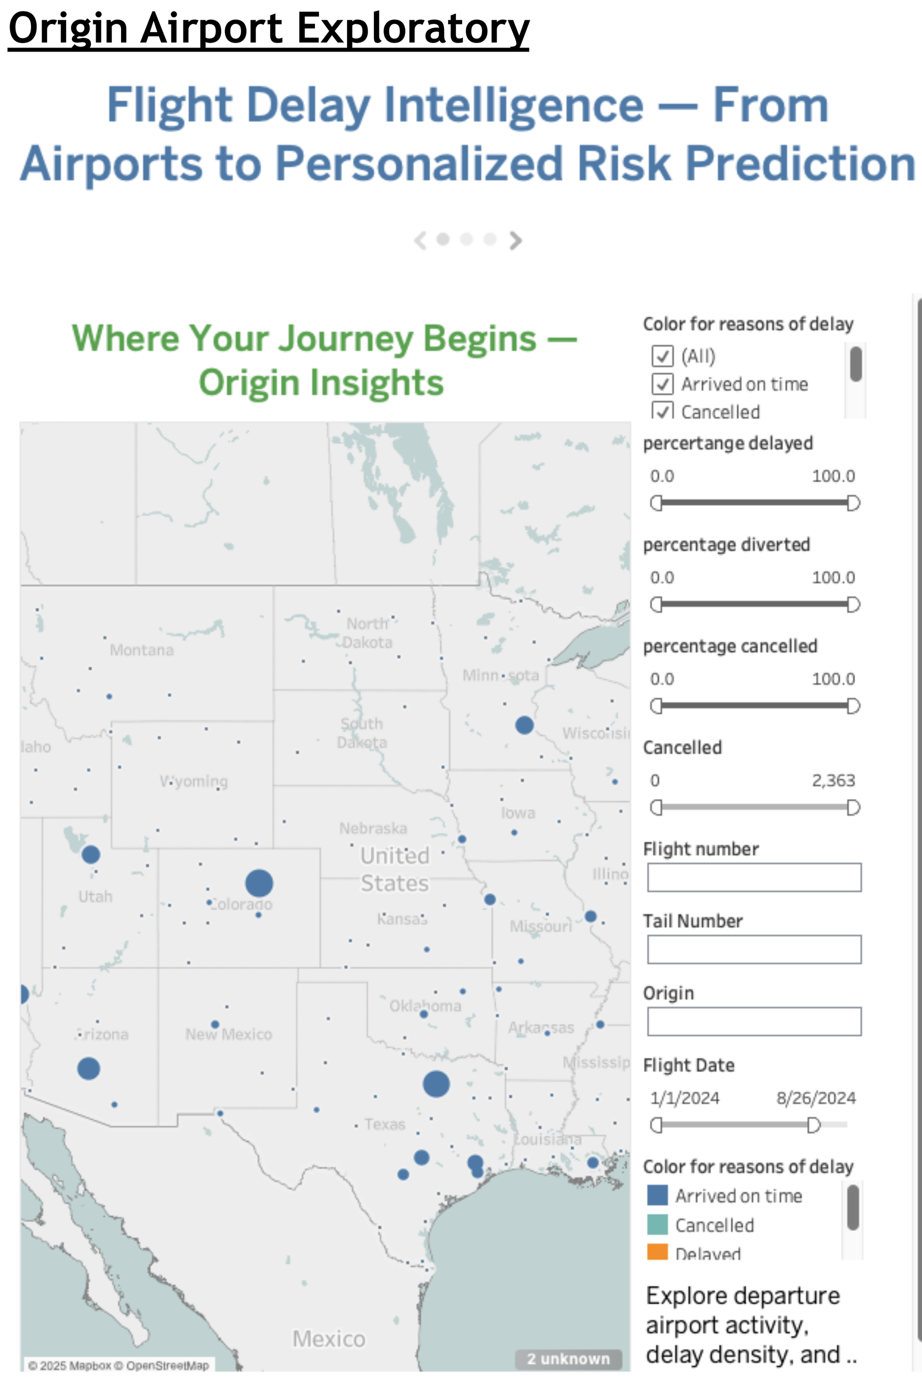

Destination Airport Exploratory

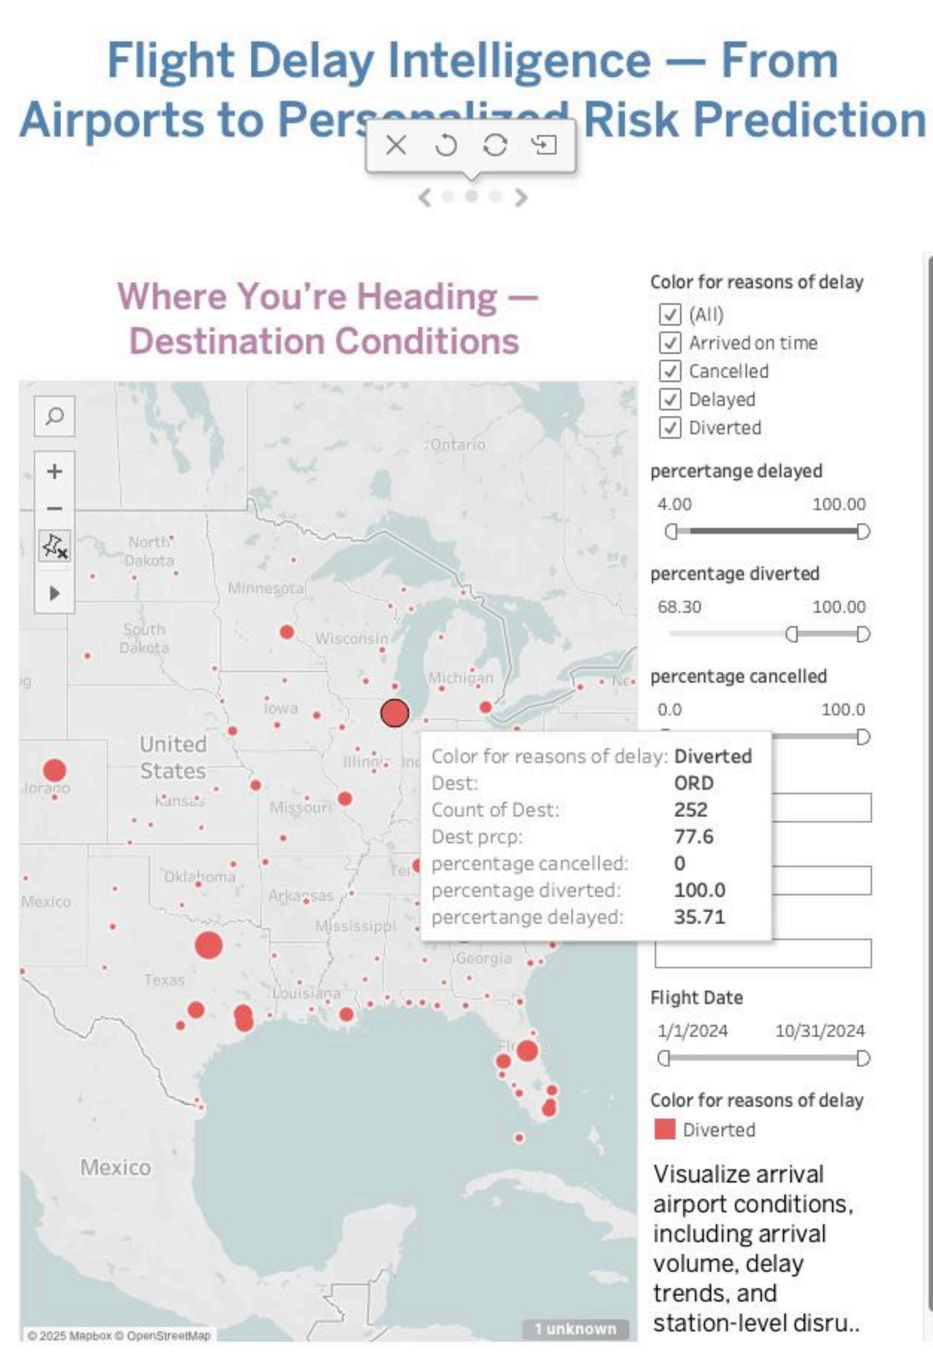

Personalized Flight Prediction