Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
*A Regional Conference of the Society for Industrial and Systems Engineering*

# ADS-B Traffic Anomaly Detection for Geopolitical Early Warning: An STL-Based Pipeline

**Eric Dailey*[1], Kristian Nordby[1], Zachary Reynolds[2], and David Beskow[2]**

[1]Department of Mathematical Sciences, United States Military Academy, West Point, New York 10996

[2]Department of Systems Engineering, United States Military Academy, West Point, New York 10996

Corresponding author's Email: ericdailey15@gmail.com

**Author Note:** The views expressed herein are those of the authors and do not reflect the position of the United States Military Academy, the Department of the Army, or the Department of Defense.

**Abstract:** Anomalies in transportation trends can serve as indicators of geopolitical disruption. This research leverages Automatic Dependent Surveillance–Broadcast (ADS-B) data from the OpenSky API to analyze aviation activity. A point-in-polygon approach aggregates daily counts of unique aircraft by country using national boundaries and ICAO24 identifiers. We evaluate two anomaly detection approaches—a hybrid Seasonal-Trend decomposition using Loess (STL) model and an Isolation Forest—on two known geopolitical events, using Iran as the positive case and Germany as a negative control. The hybrid STL introduces a parameter to detect sustained anomalous periods, which commonly arise during military interventions or other severe airspace disruptions. Across both case studies, the hybrid STL outperforms the Isolation Forest, achieving a recall of 0.91 and precision of 0.83. Such models could support early warning of geopolitical disruptions reflected in aviation patterns. Future work will focus on tuning STL parameters for robust real-time anomaly detection.

*Keywords*: ADS-B, aviation density, anomaly detection, geopolitical intelligence, early warning, time series, STL

## 1. Introduction

Global transportation patterns are direct physical manifestations of geopolitical stability. When those patterns shift abruptly—flights vanishing from a country's airspace, shipping lanes suddenly emptying—they often signal disruptions that precede or accompany major world events. Leveraging these signals as early warning indicators represents an underexplored opportunity for geopolitical intelligence.

Commercial transportation entities often receive advance notification of military interventions to mitigate civilian risk. Figure 1 illustrates this across three geopolitical conflicts: the Russian invasion of Ukraine, the U.S. intervention in Venezuela, and the U.S.-Israeli strikes on Iran. Anomaly detection applied to ADS-B aviation data can potentially surface these early warning indicators as they materialize, providing senior leaders timely signals of geopolitical disruptions affecting airspace.

This study investigates whether country-level ADS-B aircraft density anomalies can serve as early indicators of geopolitical disruptions. Using flight transponder data from the OpenSky API, we aggregate daily unique aircraft counts by country via a point-in-polygon approach, producing country-level time series as input to two anomaly detection models. We evaluate a hybrid Seasonal-Trend decomposition using Loess (STL) model—augmented with a consecutive-drop rule to capture sustained airspace suppressions—against an Isolation Forest baseline, using Iran as a positive case and Germany as a negative control. The remainder of this paper proceeds as follows: Section 2 reviews relevant literature on ADS-B anomaly detection and spatial aggregation; Section 3 describes our data collection and preprocessing pipeline; Section 4 details the STL and Isolation Forest methodologies; Section 5 presents quantitative results; and Section 6 discusses operational implications, limitations, and future work.

## 2. Background and Literature Review

ADS-B pipelines have been leveraged to detect airspace anomalies and alert authorities to unsafe conditions, with much of the existing research focused on identifying anomalous behavior at the individual aircraft level. One prominent threat involves Unmanned Aerial Systems (UAS) that emit false tracking messages to masquerade as legitimate aircraft, distorting controller situational awareness and creating significant safety risks. Pirolley et al. implemented a deep learning algorithm to

ISBN: 97819384962-5-7


(a) ADS-B flight traffic over Eastern Europe during the Russian invasion of Ukraine.

(b) ADS-B flight traffic over the Caribbean and South America during a U.S. military intervention in Venezuela.

(c) ADS-B flight traffic over the Middle East during U.S. and Israeli strikes on Iran.

Figure 1: Real-world examples of geopolitical conflict reflected in ADS-B aviation data. A notable reduction in flight density over the affected region is visible in each case, motivating the use of airspace activity as an early warning indicator of military action.

model typical flight behavior across large trajectory datasets and flag deviations indicative of UAS activity (2023). Similarly, Gariel et al. compared typical and observed trajectories to compute airspace complexity scores for air traffic controllers (2011).

While these approaches successfully classify anomalous individual trajectories, a gap remains in aggregating such anomalies to interpret broader geopolitical activity. Our model addresses this gap by applying time-series analysis to country-level aircraft density, identifying large-scale drops that may signal emerging conflicts (Saed and Omar 2025). These density drops often precede kinetic military action, as warning systems such as the FAA's Notice to Air Mission (NOTAM) system prompt airlines to avoid unsafe airspace (U.S. Department of Transportation 2023). This research therefore aims to provide early warning indicators of geopolitical events before the onset of conflict.

While we find no prior application of this approach to ADS-B data, aggregating geo-referenced point data within polygon boundaries and analyzing the resulting counts as a time series is well-established in other domains. In maritime surveillance, the Automatic Identification System (AIS)—the naval analogue to ADS-B—has been used extensively to aggregate vessel transponder locations within geographic regions and detect anomalies in traffic density; Wolsing et al. surveyed 44 such approaches, finding that AIS density anomalies reliably indicate events relevant to safety and national security (2022). In epidemiology, GPS-referenced case reports are assigned to administrative polygons and monitored as daily time series, with deviations from baseline triggering early warnings—a method the CDC identifies as standard field practice (Centers for Disease Control and Prevention 2024). We apply this same pipeline to ADS-B data, extending an established methodology to the detection of geopolitical disruptions.

## 3. Data and Preprocessing

ADS-B is a surveillance technology equipped on most commercial aircraft that transmits location and flight details to ground sensors (Federal Aviation Administration 2022). We collect ADS-B data from the OpenSky Network API via a cron job that queries current flight data every four minutes, appending observations to a daily CSV file compressed with gzip at end of day. The worldwide API coverage is depicted in Figure 2.

The OpenSky Network is an open source provider of air traffic control data that utilizes sensors across the globe furnished by local volunteers or governmental organizations. Resulting, the sensor network that it provides exhibits high coverage in the Western Hemisphere as well as Western Europe, but is sparse in Africa and most of Asia. While this is somewhat limiting, OpenSky nonetheless offers the best available high volume open source ADS-B data, and is sufficient for delivering operationally relevant insights in several regions of interest.

Fields collected from the API include timestamp, coordinates, ICAO24 identifier, velocity, and altitude. Days with incomplete data due to API interruptions or pipeline failures are filtered to ensure integrity. Figure 3 illustrates this data stream spatially, depicting aircraft density over the Persian Gulf the day before and after the onset of conflict in Iran—demonstrating the visibility of near-complete airspace shutdowns in the raw data.

Each aircraft observation is assigned to a country using a point-in-polygon operation with pure Python ray-casting against the Natural Earth 10m shapefile. Observations are de-duplicated by ICAO24 to ensure each aircraft is counted once per day, then aggregated to a single daily count per country. This produces a country-level daily time series that serves as the primary input to the anomaly detection models.

ISBN: 97819384962-5-7

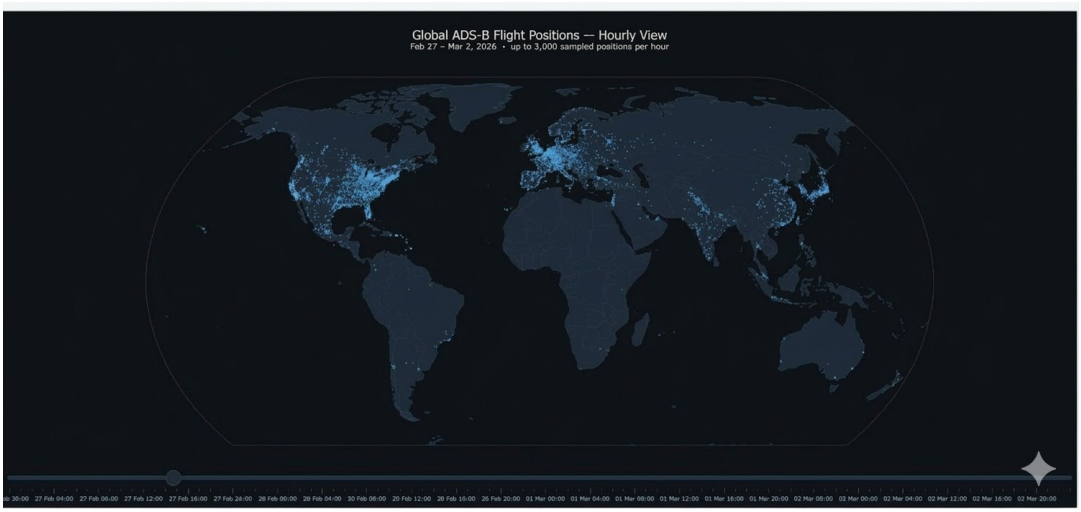


Figure 2: OpenSky API Coverage

## 4. Methodology

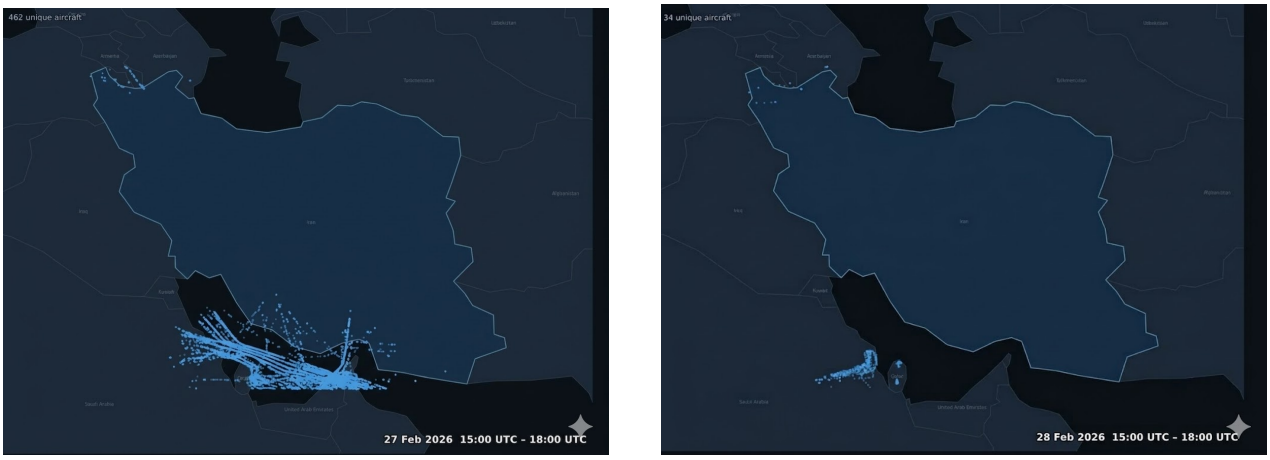


(a) Before event (27 FEB, 15:00–18:00 UTC) (b) After onset (28 FEB, 15:00-18:00 UTC local)

Figure 3: ADS-B aircraft density over the Persian Gulf before and after the detected anomaly, illustrating the near-complete airspace clearance visible in the raw data we seek to detect with through our models.

### 4.1. ADS-B Aircraft Density Anomaly Detection with STL Time Series Model

STL decomposes a time series into trend, seasonal, and residual components via locally weighted regression. We flag anomalies using residual z-scores derived from a 21-day moving average, applying a threshold of 2.5 standard deviations to balance sensitivity and specificity. A limitation of standard STL is that a persistent anomaly will cease flagging once the moving average adapts to the new baseline. To address this, we augment the model with a consecutive-drop rule: periods where aircraft counts remain below 30% of the trend for at least three consecutive days are flagged regardless of z-score, distinguishing sustained disruptions from transient dips.

The final output records, per country per day, whether the point-anomaly or consecutive-drop criterion was met, along with the observed count, STL trend value, and residual z-score. Table 1 summarizes the STL hybrid detector's key design parameters.

Table 1: STL hybrid detector design parameters.

| **Parameter** | **Value** | **Rationale** |
|---|---|---|
| Trend window | 21 days | Prevents trend absorbing sustained drops |
| Robust fitting | Enabled | Downweights outliers during decomposition |
| Point anomaly | $\|z\| > 2.5\sigma$ | Flags large single-day residuals |
| Consec.-drop threshold | 30% below trend | Captures sustained event body |
| Consec.-drop duration | $N \geq 3$ days | Filters transient one-day dips |
| Final flag | Point OR Consec. | Union detector (hybrid) |

### 4.2. Isolation Forest–Based Flight Density Anomaly Detection

The Isolation Forest (IF) takes the same input as the STL: unique aircraft per country per day. We train a separate IF per country using `contamination=0.05` and `n_estimators=200`, with a `StandardScaler` applied before fitting. Table 2 lists the seven input features. A limitation of the IF is its lack of a consecutive-drop rule, meaning sustained disruptions will cause the model to adapt rather than continue flagging anomalies.

Table 2: Isolation Forest feature set (per-country model).

| **Feature** | **Description** |
|---|---|
| unique_aircraft | Raw daily count |
| rolling_mean | 7-day centered moving average |
| rolling_std | 7-day centered standard deviation |
| pct_change_1d | Day-over-day percentage change |
| pct_change_7d | Week-over-week percentage change |
| z_from_rolling | $(y_t - \bar{y}_7)/\sigma_7$ |
| day_of_week | Cyclic weekly encoding |

## 5. Results

Figure 4 shows unique daily aircraft counts for Iran (IRN) and Germany (DEU) during the study window (4 February–11 March 2026). Iran held a stable baseline of ≈190 aircraft/day through 27 February, then collapsed to near zero beginning 28 February and remained suppressed through the end of the window. Germany maintained ≈2,700 aircraft/day with a regular weekly rhythm and a single transient dip on 8 March that recovered within one day. This contrast motivates the consecutive-drop design: Iran's suppression is *sustained*; Germany's dip is *transient*.

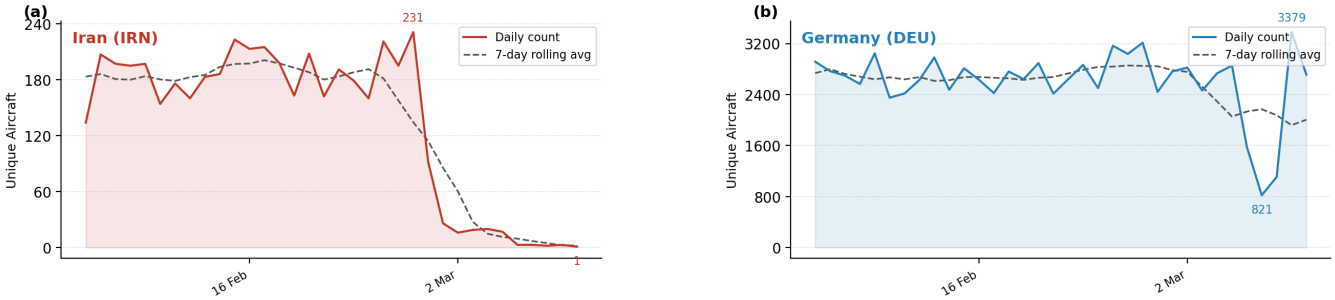


Figure 4: Daily ADS-B unique aircraft for Iran (IRN) and Germany (DEU), 4 Feb – 11 Mar 2026.

Figures 5a and 5b display the STL decomposition and Isolation Forest outputs for Iran. In the STL residual panel (Figure 5a), a pre-event spike on 27 February ($R_t \approx +125$, $z > +2.5\sigma$) triggers the point-anomaly detector one day before the collapse; sustained negative residuals from 1 March onward are captured by the consecutive-drop rule. The seasonal panel is clean (±25 aircraft/day), confirming the collapse is not a day-of-week artifact. In Figure 5b, the Isolation Forest flags $n = 2$

days at onset where the decision score drops to $\approx 0$, but the rolling mean adapts to the new low-count regime and the sustained collapse ceases to appear anomalous—the model's principal weakness for prolonged events.

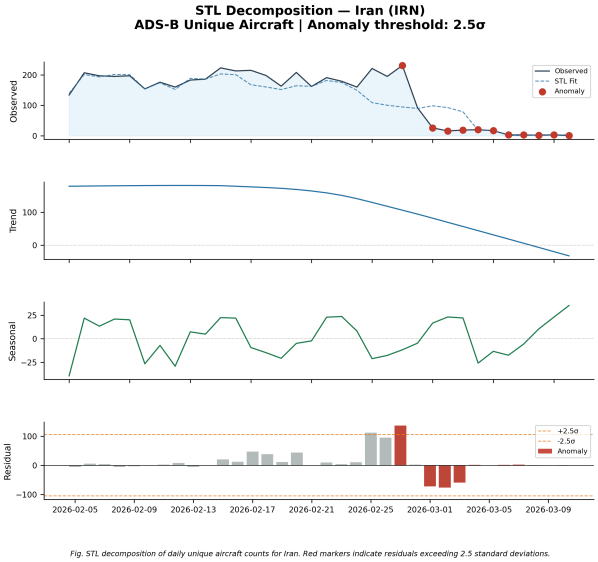


(a) STL decomposition, Iran (IRN).

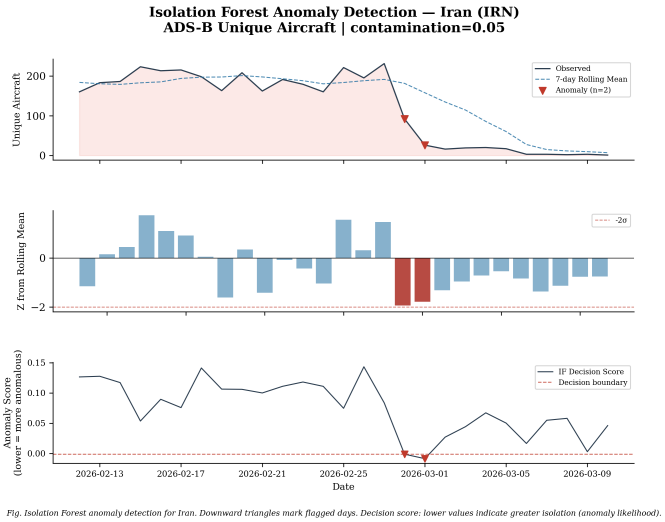


(b) Isolation Forest results, Iran (IRN).

Figure 5: Model outputs for Iran. Red markers/triangles indicate flagged anomaly days.

Table 3 summarizes quantitative performance across all model configurations. Ground truth is the Iran conflict onset period (28 Feb – 11 Mar 2026); Germany (DEU) is the negative control.

Table 3: Quantitative comparison of model configurations. Iran = positive case; Germany = negative control.

| Model | Lag | Precision | Recall | FP Rate |
|---|---|---|---|---|
| STL Point Anomaly | N/A | 0.00 | 0.00 | 0.31 |
| STL Consecutive-Drop | +1d | 0.83 | 0.91 | 0.15 |
| STL Hybrid (OR) | +1d | 0.62 | 0.91 | 0.46 |
| Isolation Forest | 0d | 0.50 | 0.18 | 0.15 |
| **STL Hybrid + IF (AND)** | **+1d** | **0.50** | **0.09** | **0.08** |

The STL consecutive-drop rule achieves the strongest recall (0.91) with reasonable precision (0.83). The Isolation Forest detects onset at $T + 0$—one day ahead of STL—but misses the sustained event body, yielding recall of only 0.18. The AND combination produces the lowest false-positive rate (0.08), favoring high-confidence alerting at the cost of recall.

## 6. Discussion

Two limitations warrant mention. First, the IF baseline adapts over a 7-day rolling window, causing anomaly scores to decay even as an event persists. A potential mitigation is re-anchoring baseline statistics to a longer historical window, preserving sensitivity to prolonged disruptions while accommodating gradual shifts in traffic patterns. Second, ADS-B sensor coverage is sparse outside the continental United States and Europe, with sensors often confined to coastal regions—as visible in Iran in Figure 3.

Operational deployment presents a precision–recall tradeoff in which recall is the more critical metric, as false negatives represent the costliest failure mode. Under this criterion, the STL Consecutive-Drop demonstrates the strongest operational potential, achieving the highest recall (0.91) and precision (0.83).

This study should be interpreted as a case study rather than a large-scale evaluation. Model performance is assessed on a single positive case (Iran) and one negative control (Germany), limiting generalizability. Future work should expand evaluation to a broader set of labeled geopolitical events across multiple regions and incorporate a higher-coverage ADS-B data source.

ISBN: 97819384962-5-7

## 7. Conclusion

This research introduces a consecutive-drop rule augmenting STL time-series decomposition to address the baseline-adaptation failure mode observed in rolling-statistic models such as the Isolation Forest. Among the benchmarked detectors, the STL Consecutive-Drop achieves the strongest balance of recall and precision, while the AND combination offers the lowest false-positive rate for high-confidence alerting at the cost of recall.

Aviation density anomalies represent an underexplored intelligence signal. Pairing the methodology introduced here with high-coverage ADS-B data could provide decision-makers a global early warning capability for emerging geopolitical conflicts.

## References

Centers for Disease Control and Prevention (2024). Geographic information system data. Accessed: 2026-03-11.

Federal Aviation Administration (2022). Automatic dependent surveillance–broadcast (ads-b). `https://www.faa.gov/air\_traffic/technology/adsb`. Accessed: 2026-03-10.

Gariel, M., Srivastava, A. N., and Feron, E. (2011). Trajectory clustering and an application to airspace monitoring. *IEEE Transactions on Intelligent Transportation Systems*, 12(4):1511–1524.

Pirolley, M., Couturier, R., Salomon, M., and Ambert, F. (2023). Ads-b anomaly detection in the surveillance of low-altitude aircrafts. *Journal of Open Aviation Science*, 1(2).

Saed, M. B. and Omar, B. (2025). Air traffic in context: Geopolitical and technical factors affecting its safety and security. https://www.researchgate.net/publication/389265035_Air_traffic_in_context_geopolitical_a nd_technical_factors _affecting_its_safety_and_security. ResearchGate.

U.S. Department of Transportation (2023). The federal aviation administration's notam system failure and its impacts on a resilient national airspace. https://www.transportation.gov/federal-aviation-administrations-notam-sys tem-failure-and-its-impacts-resilient-national-airspace.

Wolsing, K., Roepert, L., Bauer, J., and Wehrle, K. (2022). Anomaly detection in maritime ais tracks: A review of recent approaches. *Journal of Marine Science and Engineering*, 10(1):112.