Journal Pre-proof

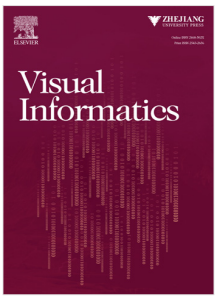

AirCascades: A visual analytics system for cascading taxiway interaction patterns

Guoqiang Wang, Xiaolin Wen, Fengjie Wang, Luge Yang,
Qipeng Wang, Min Zhu

PII: S2468-502X(26)00031-8
DOI: https://doi.org/10.1016/j.visinf.2026.100335
Reference: VISINF 100335

To appear in: *Visual Informatics*

Received date : 12 April 2026
Revised date : 9 June 2026
Accepted date : 10 June 2026

Please cite this article as: G. Wang, X. Wen, F. Wang et al., AirCascades: A visual analytics system for cascading taxiway interaction patterns. *Visual Informatics* (2026), doi: https://doi.org/10.1016/j.visinf.2026.100335.

AirCascades: A Visual Analytics System for Cascading Taxiway Interaction Patterns

Guoqiang Wang[a,b], Xiaolin Wen[c], Fengjie Wang[d], Luge Yang[a], Qipeng Wang[a] and Min Zhu[a,*]

[a]College of Computer Science, Sichuan University, Chengdu, China
[b]The Second Research Institute of CAAC, Chengdu, China
[c]Nanyang Technological University, Singapore, Singapore
[d]Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong, China

ARTICLE INFO

Keywords:
Visual Analytics
Spatiotemporal Data
Airport Operations
Interaction Cascades
Pattern Mining

ABSTRACT

Efficient management of aircraft taxiing operations depends on understanding not only where local conflicts occur, but also how they propagate through the taxiway network and evolve into cascading delays and operational risks. Existing research and operational tools still focus primarily on pairwise conflicts, raw trajectories, or coarse aggregates, making these propagation mechanisms difficult to inspect. To address this gap, we present AirCascades, a visual analytics system for post-hoc analysis of cascading taxiway interaction patterns. AirCascades uses a semantics-driven abstraction pipeline that transforms trajectories into encounter sessions, links them into cascade instances through shared-aircraft continuity, resource overlap, and topology-aware coupling cues, and clusters these instances into motifs that expose recurring propagation structures. Building on this hierarchy, the interface provides a macro-level view for system-wide screening, a meso-level pattern view with map-free glyphs for structural comparison, and a micro-level replay view for tracing individual cascades through speed profiles and inter-event links. A coupling-sensitivity analysis, a structural audit of the default analysis set, and qualitative expert walkthroughs on 24 hours of operational data from a major international airport suggest that AirCascades helps analysts inspect propagation mechanisms and recurring structural bottlenecks that conventional surveillance replay does not readily expose.

1. Introduction

After each operational day, airport surface managers review conflicts and accumulated delays. In this paper, a *conflict* is a taxiway encounter in which aircraft proximity, yielding, stopping, or speed reduction indicates reduced separation margin or efficiency loss. A harder question is whether one local interaction triggered later encounters at other taxiway locations. For example, unexpected braking at a runway holding position can slow following aircraft, build queues at upstream intersections, and ripple through the network. Prior studies show that surface congestion at mega-airports has cascade-type effects where a single delay accumulates over subsequent hours (Jiang, Xue, Wang, Cai, Zhang and Li, 2021), and that network-wide delay propagation can involve large parts of the airport network (Fleurquin, Ramasco and Eguiluz, 2013). Understanding these mechanisms supports targeted operational changes, controller training, and safety reviews.

Current approaches for airport surface management focus mainly on pairwise conflict detection or static hotspot identification (locations with historically high conflict frequencies) (Kuchar and Yang, 2000; Vaddi, Kwan, Fong and Cheng, 2012; Albrecht, Lee and Pang, 2012). While effective for raising immediate alerts, these methods treat each encounter as an isolated incident (Figure 1(a)). In our terminology, such an encounter is represented as a **Spatio-temporal Encounter Session (SES)**, a behavioral record of aircraft interactions near a taxiway interaction point. Existing tools offer little support for tracing how separate SES events are connected. When a conflict at one location leads to subsequent interactions elsewhere, operations analysts still lack systematic support for following the propagation chain or recognizing recurring patterns that could inform procedure optimization.

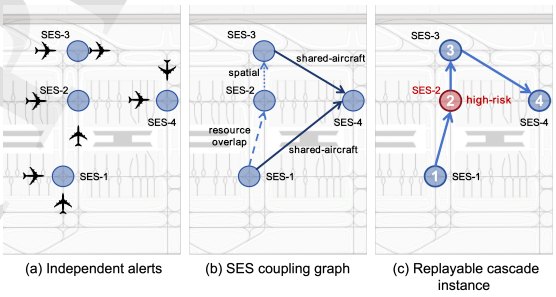

Figure 1: From isolated alerts to replayable cascade instances. (a) Conventional surface-conflict analysis exposes Spatio-temporal Encounter Session (SES) events as separate alerts. (b) AirCascades organizes the same SES events into an evidence-coupled graph with shared-aircraft, resource-overlap, and spatial-reachability links. (c) The system then materializes one retained path-instance for replay and analysis, with the highlighted node indicating elevated interaction severity.

AirCascades instead organizes the same SES events into an evidence-coupled graph (Figure 1(b)) and extracts retained path-instances for replay (Figure 1(c)). Here, *evidence-coupled* means that links are supported by operational

*Corresponding authors
xiaolin004@e.ntu.edu.sg (X. Wen); zhumin@scu.edu.cn (M. Zhu)

continuity cues rather than by a single linkage rule. These cues include **shared-aircraft continuity**, where the same aircraft appears in adjacent SES events; localized resource overlap, where encounters compete for the same node or segment; and topology-aware reachability, where a later event is feasible along the taxiway network. Related work in visual analytics suggests that such propagation structures become analyzable when events are embedded in spatial or topological context. Systems for traffic congestion (Wang, Lu, Yuan, Zhang and van de Wetering, 2013) and air pollution propagation (Deng, Weng, Chen, Liu, Wang, Bao, Zheng and Wu, 2020) demonstrate that pattern mining and topology-aware visualization can reveal how disturbances spread through spatial networks. However, these approaches assume fundamentally different propagation mechanisms, such as backward-moving traffic waves or wind-driven pollutant transport. Airport surface operations are governed by strict separation rules and explicit control commands (ICAO, 2016). Propagation occurs through discrete operational couplings: shared-aircraft continuity, sequential dependencies at merge points, and capacity constraints at runway thresholds. This combination distinguishes the domain from existing cascade models and calls for purpose-built abstractions.

To ground our work in real operational needs, we conducted a four-month formative study with three domain experts; their backgrounds and roles in the iterative design process are detailed in Section 3.3. Through field observations, semi-structured interviews, and prototype reviews, we identified a systematic gap in existing tooling. Post-operations analysis is routine but labor-intensive: analysts must manually correlate many conflict alerts across time and space, a process the experts described as "assembling puzzle pieces without a picture of the box." The key finding from our formative study is that surface operations analysis is fundamentally a *post-hoc* activity. This workflow shaped the main design emphasis of AirCascades: ranking cascades first, then supporting spatial localization, pattern comparison, and case-level replay.

Our collaboration revealed three challenges that existing tools do not adequately address:

- **C1: From trajectories to meaningful interaction events.** Raw surveillance data captures aircraft positions, but interaction analysis requires identifying *where* and *when* aircraft constrain each other and summarizing each encounter with attributes such as severity and delay.
- **C2: From isolated events to propagation chains.** Local interactions become consequential when they propagate. Linking individual encounters into cascades that capture "who affected whom, where, and when" requires a model that preserves explicit shared-aircraft continuity when available while accounting for temporal, resource, and network constraints.
- **C3: From individual cascades to recurring patterns.** Airport operations generate thousands of cascades daily. Experts need tools to discover motif families of structurally similar cascades and compare their topological signatures, risk profiles, and temporal contexts without geographic clutter.

We address these challenges with **AirCascades**, a visual analytics system for identifying, summarizing, and analyzing cascading taxiway interaction patterns (Figure 1(c)). Our approach has two main components. First, we introduce a *semantics-driven modeling pipeline* (C1–C2) that transforms raw trajectories into four derived levels: **Geometric Interaction Points (GIPs)** where interactions can occur, SES events with behavioral descriptors, cascades that link SES through evidence-weighted operational couplings, and motifs that group cascades by structural similarity. Two cascades occurring at different parts of the airport are therefore treated as structurally similar if they share the same canonical interaction progression, such as an H-X-M sequence (head-on → crossing → merge), even when their geographic footprints differ.

Second, we design a *three-level visual analytics framework* (C3). The macro-level overview uses a dual-layer map to separate network topology from propagation patterns and helps analysts identify cascade-prone regions. The meso-level pattern view uses topology-encoding glyphs that abstract away location, allowing experts to compare cascade families by structural signatures. The micro-level narrative view uses synchronized speed-curve cards and shared-aircraft connection arcs to support frame-by-frame replay of cascade dynamics. These views are coordinated through brushing and linking: selecting a cascade in the leaderboard highlights its spatial footprint and populates the micro-level replay, while clicking a pattern glyph filters the map to show matching instances.

We evaluate AirCascades on 24 hours of real operational data from Beijing Capital International Airport (PEK). A coupling-sensitivity analysis, a structural audit of the default 104-cascade set, and qualitative expert walkthroughs indicate that the system helps analysts inspect propagation mechanisms and recurring structural bottlenecks that conventional surveillance replay does not readily expose.

Together, the paper contributes:

1. **An evidence-coupled cascade modeling pipeline** that turns Automatic Dependent Surveillance-Broadcast (ADS-B) trajectories into SES events and replayable cascade instances using shared-aircraft continuity, resource overlap, and topology-aware reachability. The pipeline avoids a single linkage rule and enables motif mining at scale.
2. **AirCascades, a coordinated macro-to-micro visual analytics interface** for cascade triage, spatial localization, structure-first pattern comparison, and micro-level narrative replay that separates shared-aircraft anchors from weaker time- and topology-based continuity cues.

3. **An evaluation on a major hub dataset** that tests interpretability under different coupling modes and examines whether the interface supports domain practitioners' post-hoc analysis. Coupling sensitivity, structural audits, and qualitative expert walkthroughs indicate a practical balance among cascade coverage, evidential strength, and analytical utility.

2. Related Work

Our work draws on five research areas: airport surface conflict modeling, air traffic and airport surface visualization, visual analytics for trajectory and mobility data, cascading and propagation pattern analysis, and graph motif mining and network visualization. We review each area briefly and position AirCascades relative to the closest precedents.

2.1. Airport Surface Conflict Detection and Risk Modeling

Conflict detection and resolution (CDR) has been extensively studied in air traffic management. Kuchar and Yang's taxonomy (Kuchar and Yang, 2000) established foundational concepts for en-route operations, but focused on pairwise conflict geometry rather than constrained surface networks. For airport surfaces, researchers have developed automated detection and risk assessment methods: MAESTRO (Vaddi et al., 2012) combines trajectory prediction with separation criteria to detect potential conflicts, while Albrecht et al. (Albrecht et al., 2012) visualize conflict probabilities to reveal high-risk regions.

More recent work addresses surface conflict analysis with machine learning, simulation, and hotspot assessment methods. Spatio-temporal overlap analysis identifies persistent hotspot locations (Xia, Zheng, Wan and Zhu, 2019; Ali, Delair, Pham, Alam and Schultz, 2020), Monte Carlo-based risk modeling identifies key surface hotspots under mixed traffic flows (Tian, Zhou, Yin, Li and Zhang, 2024), and data-driven recognition of hot taxiway segments extends hotspot analysis from conflict points to operationally busy segments (Yin, He, Ma, Qiao and Liu, 2025). Graph neural networks predict surface trajectories and infer future conflict risks (Zhang, Zhong and Mahadevan, 2022), and control command inference enables early warning of potential separation violations (Zhuge, Liang, Zhang, Li, Yang and Wu, 2024). ADS-B-based surface studies further extract taxi routes, runway usage, taxi times, and operational performance indicators for optimization at PEK (Ma, Zhou, Liang and Delahaye, 2022). Cellular automaton models have also been used to study congestion formation at mega-airport surfaces, demonstrating empirically that a single interaction event can propagate into broader system effects (Jiang et al., 2021). At the network level, Fleurquin et al. (Fleurquin et al., 2013) show that delay propagation in airport networks is systemic and network-wide, which motivates the cascade framing adopted here.

Despite this breadth, existing methods still analyze surface risk primarily as isolated events, hotspots, or predicted conflicts. A detected conflict at intersection A is evaluated without systematic support for tracing whether it influenced aircraft at intersection B, or for recognizing recurring structural patterns across such chains. AirCascades addresses precisely this gap.

2.2. Air Traffic and Airport Surface Visualization

While Section 2.1 covers algorithmic conflict detection, a separate body of work addresses the *visualization* of air traffic and airport surface operations. At the en-route scale, trajectory bundling techniques (Hurter, Ersoy and Telea, 2013) reduce visual clutter in dense flight-path displays by clustering spatially proximate trajectories into smooth aggregated representations. This design challenge is structurally related to our cascade path visualization, but occurs at a much larger spatial scale and without propagation semantics. Andrienko et al. (Andrienko and Andrienko, 2013; Andrienko, Andrienko, Bak, Keim and Wrobel, 2013) provide broader frameworks for movement pattern analysis, including trajectory clustering, temporal segmentation, and interactive querying, that inform our treatment of SES events as structured behavioral records.

Airport surface operations are governed by Advanced Surface Movement Guidance and Control Systems (A-SMGCS) (ICAO, 2004), whose visual interfaces present raw positional data as moving labels on an airport diagram. These systems provide limited abstraction above the individual aircraft level and limited support for post-hoc analysis of historical interaction patterns. Albrecht et al. (Albrecht et al., 2012) extend surface displays with probability-based conflict visualization, but the representation remains event-centric and does not readily express propagation relationships between successive encounters.

Recent visual analytics interfaces for air traffic management also emphasize real-time decision support. Zohrevandi et al. (Zohrevandi, Vrotsou, Westin, Lundberg and Ynnerman, 2024) design a real-time visual analytics interface for managing air traffic complexity through live traffic-state inspection and scenario comparison. This is close to our operational domain, but it still targets online complexity management rather than retrospective airport-surface cascade analysis. SeRViz (Jalilvand, Milhomem and Paulovich, 2025) shows the value of interactive visual analytics for airport ground handling sequences, but its focus is sequential rule exploration over ground-service logs rather than taxiway trajectory propagation.

The gap is therefore not a lack of airport visualization per se, but the absence of post-hoc, cascade-aware visual analytics for airport surface operations. Unlike en-route systems that aggregate trajectories at scale, AirCascades must preserve individual encounter semantics; unlike A-SMGCS displays aimed at live monitoring, it is designed for exploratory retrospective analysis where the cascade chain itself becomes the analytical unit.

2.3. Visual Analytics for Trajectory and Mobility Data

The visual analytics community has developed rich methods for spatiotemporal movement data. Surveys by Deng et al. (Deng, Weng, Liu, Tian, Xu and Wu, 2023), Andrienko et al. (Andrienko, Andrienko, Chen, Maciejewski and Zhao, 2017), and Chen et al. (Chen, Guo and Wang, 2015) emphasize progressive abstraction from raw trajectories to semantically meaningful patterns, together with the integration of human expertise through interactive exploration. Andrienko and Andrienko (Andrienko and Andrienko, 2013) further outline movement-oriented analysis methods covering clustering, pattern mining, and interactive querying that inform our design.

Urban-scale trajectory systems address movement largely at aggregate levels. Ferreira et al. (Ferreira, Poco, Vo, Freire and Silva, 2013) pioneered visual exploration of massive taxi trip data with interactive temporal queries; subsequent work studies route choice behavior (Lu, Lai, Ye, Liang and Yuan, 2017), semantic trajectory interaction (Al-Dohuki, Kamw, Zhao, Yang, Li, Wu, Ye, Chen, Ma and Wang, 2017), traffic flow density visualization (Scheepens, Willems, van de Wetering, Andrienko, Andrienko and van Wijk, 2011), and congestion cause analysis (Pi, Yeon, Son and Jang, 2021). RoseTrajVis (Ferreira, Afonso, Ferreira and Vaz, 2022) uses movable rose-diagram lenses to summarize local trajectory directions and speed bands, while HoLens (Feng, Zhu, Wang, Hao, Yang, Zeng and Qu, 2024) models higher-order movement patterns through adaptive aggregation and state-sequence visualization. Trajectory similarity research has also matured substantially, with recent surveys systematizing distance measures and quantitative comparisons for spatio-temporal trajectories (Hu, Chen, Fang, Fang, Li and Gao, 2024). At the microscopic level, TripVista (Guo, Wang, Yu, Zhao and Yuan, 2011) analyzes intersection traffic from multiple perspectives, while Wang et al. (Wang, Chen, Wu, Zhao, Hong, Gu, Wang, Liang and Bao, 2014) support data-driven transport assessment on urban roads. MobilityGraphs (von Landesberger, Brodkorb, Roskosch, Andrienko, Andrienko and Kerren, 2016) shows how spatio-temporal abstraction over graph structures can reveal mass mobility dynamics, and Zeng et al. (Zeng, Fu, Müller Arisona, Erath and Qu, 2016, 2014) visualize origin-destination patterns and public transportation flows, concepts we adapt for cascade path visualization.

Visual analytics for spatial event pattern analysis provides complementary foundations. Von Landesberger et al. (von Landesberger, Bremm, Andrienko, Andrienko and Tekušová, 2012) show how interactive visual analytics can reveal structured patterns in categoric spatio-temporal data through aggregation, comparison, and coordinated exploration, an approach echoed in our motif-level exploration. Maciejewski et al. (Maciejewski, Rudolph, Hafen, Abusalah, Yakout, Ouzzani, Cleveland, Grannis and Ebert, 2010) develop predictive visual analytics methods for identifying spatiotemporal hotspots from historical data, conceptually related to our ISI-ranked cascade leaderboard for directing analyst attention toward high-severity incidents.

These systems contribute design principles we incorporate, including multi-view coordination (Roberts, 2007), task-driven abstraction, linked spatial-temporal-attribute exploration, and structure-aware comparison of movement patterns. However, urban mobility systems typically operate at aggregate levels, such as traffic volumes on road segments, OD flows between zones, or similarity groups among trajectories, rather than capturing individual interaction events with domain-specific semantics. Airport taxiing demands finer granularity: we must identify not just that two aircraft were spatially proximate, but *how* they interacted, *who yielded to whom*, and *what behavioral signatures* characterized the encounter. Our SES abstraction addresses this need by defining operationally meaningful interaction events at the GIP level.

2.4. Cascading and Propagation Pattern Analysis

The closest methodological precedents for AirCascades come from systems that explicitly model propagation processes in spatiotemporal domains. VisCas (Deng, Weng, Liang, Bao, Zheng, Schreck, Xu and Wu, 2022) presents a general visual cascade analytics framework for spatiotemporal events, supporting inference and interactive exploration of cascading patterns across urban contexts and demonstrating that cascade analysis is a valuable visual analytics paradigm beyond domain-specific applications.

Wang et al. (Wang et al., 2013) analyze traffic congestion propagation by mapping taxi trajectories to road networks, detecting low-speed segments, and linking temporally adjacent congestion events into propagation trees. Their core insight is that cascading phenomena can be captured as directed graphs connecting localized events through a shared spatial medium, and this directly inspires our cascade construction. However, traffic congestion propagates via backward-moving density waves where no individual vehicle serves as the explicit carrier of delay; the mechanism is fluid and aggregate. In the taxiway domain, propagation is discrete and operationally constrained: explicit shared-aircraft continuity provides strong evidence for linking encounters, but analysts also need topology-aware cues such as resource overlap and short-range reachability for system-wide coverage. This motivates our SES abstraction, evidence-weighted coupling formulation, and coupling-sensitivity evaluation.

Deng et al. (Deng et al., 2020) propose AirVis for air pollution propagation, modeling pollutant transport as graphs and applying frequent subgraph mining to extract recurring patterns. Their use of hierarchical motif organization and pattern-comparison visualization directly informs our motif view design. The key structural difference lies in the propagation medium: pollutant transport follows continuous atmospheric diffusion at regional scales, whereas taxiway cascades propagate through point-level aircraft interactions within a deterministic routing network. AirVis's graph nodes represent city-scale regions connected by wind fields; our

GIP nodes represent taxiway points linked by aircraft whose identity, speed profile, and compliance behavior are observable. This actor-centric model supports behavioral characterization of each cascade link, which domain-agnostic propagation frameworks usually lack.

Related work on co-occurrence pattern detection (Wu, Weng, Deng, Bao, Xu, Wang, Zheng, Ding and Chen, 2021; Li, Chen, Zhang, Andrienko and Andrienko, 2019) provides techniques for identifying spatiotemporal event correlations, but does not model the directional propagation semantics used here. Together, these precedents show the value of propagation-aware visual analytics and the need for airport-specific cascade models.

2.5. Graph Motif Analysis and Network Visualization

Recurring structural patterns in graphs, often described as motifs, are a powerful lens for understanding complex networks. Von Landesberger et al. (von Landesberger, Kuijper, Schreck, Kohlhammer, van Wijk, Fekete and Felkel, 2011) survey visual analysis methods for large graphs, covering clustering, pattern extraction, and interactive exploration of graph structure. Their taxonomy of motif types and encoding strategies informs the design of our structural Motif Glyph. MobilityGraphs (von Landesberger et al., 2016) extends this to dynamic spatio-temporal graphs, demonstrating that graph abstraction can reveal recurrent patterns invisible in raw movement data.

Geographic network visualization has also been extensively studied. Bach et al. (Bach, Dragicevic, Archambault, Hurter and Carpendale, 2017) review space-time cube approaches for temporal geographic data and identify design principles for representing both spatial structure and temporal evolution in network-embedded datasets. Their discussion of the trade-off between spatial faithfulness and temporal clarity motivates our dual-layer map design, which separates taxiway topology from propagation paths.

Recent graph-visualization work also reinforces the value of motif-based abstraction for large networks. AdaMotif (Zhou, Lai, Sun, Chen, Chen, Wu and Wang, 2025) replaces structurally similar subgraphs with adaptive motifs to preserve global structure while reducing visual complexity. This is methodologically adjacent to our motif abstraction, but its goal is generic large-graph simplification rather than identifying recurring propagation configurations in a spatiotemporal operational network.

Temporal event sequence analysis provides tools for identifying recurring patterns in ordered event data. Monroe et al. (Monroe, Lan, Lee, Plaisant and Shneiderman, 2013) show how temporal event sequences can be simplified and clustered to reveal recurring structure, an idea we adapt in our cascade fingerprinting pipeline, where each cascade is encoded as an ordered feature vector for motif clustering. The key adaptation is that cascade sequences are not flat event lists but directed propagation chains in which shared-aircraft continuity provides an operational anchor beyond temporal co-occurrence.

Despite substantial progress in graph motif analysis and geographic network visualization, existing methods target general-purpose graphs or aggregate spatial networks and are less suited to recurring *cascade propagation patterns* on operationally constrained directed networks where identifiable actors mediate propagation. Our motif abstraction addresses this gap by combining canonical interaction signatures with behavioral attributes (ISI, timing, aircraft roles) into features tailored to taxiway analysis.

2.6. Positioning Summary

AirCascades occupies a distinctive position at the intersection of these five research areas. It draws operational semantics and separation constraints from airport surface conflict analysis, dense spatial context from air traffic visualization, multi-level exploration strategies from mobility visual analytics, directed propagation modeling from cascade analysis, and structural comparison techniques from graph pattern mining. The unifying contribution is a domain-specific pipeline that turns these ingredients into a cohesive visual analytics system for *spatially constrained interaction cascade analysis*; to our knowledge, no prior system addresses this specific problem formulation.

3. Domain Background and Problem Formulation

This section motivates AirCascades, defines its core abstractions, and derives the analysis tasks that shape the system design.

3.1. Airport Surface Operations

Large airports operate under tight capacity and safety constraints. During peak periods, many aircraft occupy runways, taxiways, and aprons. Although surface movement guidance systems provide surveillance and routing support, controllers must still anticipate how local disturbances will propagate through the network (Vaddi et al., 2012).

Aircraft movements are fundamentally interdependent. A departure queue can back up from a runway entry, aircraft compete for gaps at busy intersections, and wake turbulence or runway occupancy rules impose additional temporal constraints. When one aircraft slows unexpectedly after missing a gap or receiving a late clearance, following aircraft brake, and the resulting perturbation can ripple through multiple crossings before dissipating. Such *cascading interactions* affect both safety (reduced separation margins) and efficiency (extended taxi times), yet current tools focus on isolated conflicts or static hotspots rather than tracing how disturbances propagate through the surface network.

3.2. Data Abstractions

Our analysis requires progressively abstracting raw trajectory data into interaction-centric representations. Such multi-level abstraction is a standard strategy in visual analytics for scaling from raw observations to interpretable analytical units (Andrienko and Andrienko, 2013; Sedlmair, Meyer and Munzner, 2012). We therefore introduce four

levels of abstraction: **Geometric Interaction Points (GIPs)** identify where interactions can occur, **Spatio-temporal Encounter Sessions (SES)** capture what happens when aircraft meet, **cascades** link related SES into propagation chains, and **motifs** group cascades by structural similarity.

3.2.1. Geometric Interaction Points (GIPs)

Not all taxiway locations matter equally for interaction analysis. We identify **Geometric Interaction Points** as locations where aircraft movements can constrain each other: intersections where paths cross, merge points where traffic combines, and segments or apron areas where proximity forces speed adjustment. Each GIP has a spatial footprint and one label from a four-type vocabulary: **crossing (X)**, **merge (M)**, **head-on (H)**, or **proximity (P)**. The proximity label covers same-segment following and apron near-encounters.

3.2.2. Spatio-temporal Encounter Sessions (SES)

A **Spatio-temporal Encounter Session** captures what happens when aircraft actually interact near a GIP. Rather than simply noting that two aircraft were proximate, an SES records the behavioral signatures of the encounter: which aircraft arrived first, whether either slowed or stopped, how close they came, and how long the interaction lasted. We compute an **Interaction Severity Index (ISI)** that combines risk exposure (proximity, closure rate) with efficiency loss (speed reduction, delay) into a single score for ranking and exploration across SES types.

3.2.3. Cascades

Individual SES become operationally significant when they connect into propagation chains. We construct **cascades** by linking SES through operational couplings that are feasible in time and on the taxiway network: explicit shared-aircraft continuity provides the strongest evidence, while resource-overlap and short-range network reachability provide weaker but still operationally meaningful links when direct aircraft continuity is absent. At the coupling level, these relationships form a directed acyclic graph. For visual analysis, however, AirCascades materializes each retained root-to-leaf path as a **cascade instance**, producing a linear narrative for ranking, comparison, and replay.

3.2.4. Motifs

Real operations generate thousands of cascades, many sharing similar structures. We cluster cascade instances by their canonical interaction-type signature and summary behavioral attributes into **motifs**, which are families representing recurring propagation patterns. A motif might capture "runway-queue spillback affecting an upstream intersection" or "multi-way blockage radiating from a central merge." Motifs bridge raw events and pattern-level insights for planning.

3.3. Formative Study and Analysis Tasks

Following the design study methodology proposed by Sedlmair et al. (Sedlmair et al., 2012), we collaborated with three domain experts over a four-month period to understand operational needs and iteratively refine our data abstractions. Leveraging our institutional affiliation with an aviation research organization, we had regular access to airport operations personnel and could conduct on-site observations of real-world surface traffic management.

The experts were selected to cover complementary operational perspectives. **E1** is a surface movement controller at Beijing Capital International Airport (PEK) with eight years of experience and routinely reviews daily conflict logs for controller debriefing and training. **E2** is an airport-operations expert from another major airport with 12 years of traffic-flow management experience and focuses on recurring congestion patterns and procedure optimization. **E3** is an aviation-systems specialist with six years of experience supporting safety-analysis workflows and reconstructing incident sequences for safety management reports.

The collaboration proceeded in three stages. First, through field observations and bi-weekly discussions, the experts explained airport surface operations, typical conflict scenarios, and limitations of existing analysis tools. Second, we presented initial data abstractions (GIP, SES) and prototype visualizations to test whether our representations aligned with operational reasoning. Third, we iteratively refined the cascade construction rules and visual encodings based on expert critiques, including adjustments to ISI weighting to better reflect domain priorities.

Through this process, we identified three core analysis needs that shape the AirCascades design:

T1: System-wide situational awareness. Where and when do cascading interactions concentrate? Which network structures contribute most to propagation? Experts need an overview that localizes cascade-prone regions and highlights critical structures across the entire surface.

T2: Pattern-level characterization. What recurring propagation structures exist, and how do they differ in severity, frequency, and context? Experts need tools to compare motif families without geographic clutter.

T3: Instance-level explanation. How did a specific cascade unfold? What timing relationships or decisions preceded it? For diagnosis and training, experts need to reconstruct individual cascades and examine their internal dynamics.

These tasks map directly to the three-level visual analytics framework described in Section 5: a macro-level overview for T1, a meso-level motif view for T2, and a micro-level narrative view for T3.

3.4. Design Requirements

Grounded in the analysis tasks and expert feedback from prototype walkthroughs, we derived five design requirements. Each requirement maps to one or more analysis tasks and design goals (Section 3.5).

R1: Prioritized cascade triage. Analysts need to quickly identify the most severe or complex cascades from a full operational day before investing time in detailed examination. The system must rank cascades by composite impact, exposing key attributes such as event count, severity trajectory, and

direct shared-aircraft anchors without requiring drill-down for initial screening. *(from T1)*

R2: Spatial localization within network context. Individual cascades must be positioned on the taxiway network so analysts can identify which intersections and routes are involved and how the propagation path traverses the topology. Network structure and propagation dynamics must be visually separable to avoid conflation of the physical substrate with the cascade event. *(from T1, T2)*

R3: Structure-first pattern comparison. When comparing cascade families, geographic coordinates are a distraction: two structurally identical cascades occurring at different airport locations should be perceived as similar. The system must abstract each cascade pattern into a representation that exposes topological structure and behavioral profile while suppressing positional variation. *(from T2)*

R4: Sequential interaction replay at the event level. Mechanism tracing and training preparation require reconstructing the exact timing, behavioral signatures, and direct aircraft continuities within individual cascades. The system must support fine-grained, step-by-step playback that makes explicit shared-aircraft links visible when they exist, while preserving full-cascade temporal context. *(from T3)*

R5: Unified severity quantification. Cascades differ in both safety risk (how close aircraft came) and efficiency loss (how much each aircraft decelerated). Experts need a single comparable metric that integrates both dimensions so they can rank events, filter by severity threshold, and compare instances across different cascade families. *(from T1, T2, T3)*

3.5. Design Goals

We translate the five requirements into six concrete design goals that directly govern the system components described in Section 5. Requirements and goals are cross-referenced to support traceability from domain needs to implementation decisions.

G1: Cascade Leaderboard with embedded visual encodings. *(from R1)* Design a sortable, filterable ranked list (View A) in which proportional bars, risk-trajectory sparklines, and shared-aircraft markers are embedded directly into table cells. Analysts must be able to triage the full day's cascades through visual scanning alone, without opening details.

G2: Dual-layer geospatial map. *(from R2)* Design a map view (View B) that architecturally separates the static taxiway network (rendered as a reference layer) from the dynamic propagation path (rendered as a Bezier-curve overlay). This separation must allow analysts to perceive *where the cascade occurred* and *how it traversed the network* as distinct analytical questions.

G3: Map-free topology-encoding motif glyphs. *(from R3)* Design a glyph representation (View C) that encodes the canonical interaction signature and structural profile of each cascade family without geographic coordinates. Glyphs must support side-by-side structural comparison across motif families and direct click-through filtering to the leaderboard and map.

G4: Synchronized micro-replay with shared-aircraft arcs. *(from R4)* Design a micro-level narrative view (View D) in which each SES event is represented as a speed-curve card, and shared-aircraft propagation links are rendered as explicit arcs connecting adjacent cards whenever such direct continuity exists. The replay must support frame-by-frame animation with a synchronized cursor and coordinated map display.

G5: Interaction Severity Index (ISI) as a first-class signal. *(from R5)* Define and implement a composite ISI metric that integrates proximity-based risk exposure with speed-reduction-based efficiency loss. The ISI must be surfaced consistently across all views, including as a sorting key in the leaderboard, a color dimension in glyphs and sparklines, and a header label on speed-curve cards, so that severity serves as a unified analytical language throughout the system.

G6: Cross-level coordinated interaction. *(from R1-R4)* All views must be tightly linked through brushing and linking so that selections and filters propagate across abstraction levels. An analyst should be able to move fluidly from overview triage (G1) through spatial localization (G2) and pattern comparison (G3) to instance-level replay (G4) without losing context or repeating navigation steps.

3.6. Terminology Reference

For readability, we use the following terms consistently throughout the paper. A **Geometric Interaction Point (GIP)** is a location on the taxiway network where aircraft movements can constrain each other, labeled as crossing (X), merge (M), head-on (H), or proximity (P). A **Spatio-temporal Encounter Session (SES)** is a behavioral record of an actual interaction between two or more aircraft near a GIP, including arrival order, speed changes, stops, and minimum separation.

A **cascade** is the linear propagation path extracted from the underlying SES coupling DAG and used as the primary analysis unit in the interface. Links may be supported by **shared-aircraft continuity** (the same aircraft participating in adjacent SES), resource overlap, or spatial reachability, and direct shared-aircraft links are surfaced explicitly in the interface when present. A **motif** is a cluster of structurally similar cascade instances, identified from their canonical interaction signatures and summary behavioral features. **Interaction Severity Index (ISI)** denotes the composite metric that quantifies the combined risk exposure (proximity, closure rate) and efficiency loss (speed reduction, delay) of an SES encounter.

4. Semantics-driven Modeling Pipeline

This section transforms raw trajectory data into the SES–cascade–motif hierarchy used by AirCascades. Figure 2 summarizes the four-stage pipeline.

4.1. Map Matching and Network Representation

Automatic Dependent Surveillance-Broadcast (ADS-B) is a surveillance data source in which aircraft periodically

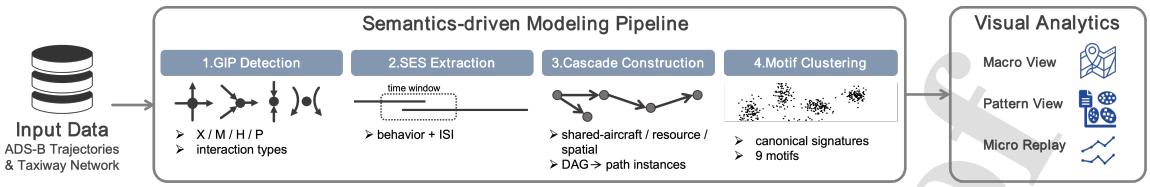

Figure 2: Overview of the AirCascades pipeline. ADS-B trajectories and the taxiway network are transformed through four stages: GIP detection identifies interaction types, SES extraction summarizes encounter behavior and computes ISI, cascade construction links SES into evidence-coupled DAGs and retained path-instances, and motif clustering groups canonical cascade signatures into recurring families. These abstractions feed the Macro View, Pattern View, and Micro Replay in the visual analytics interface.

broadcast state information. In our dataset, each record includes a timestamp, aircraft identity (e.g., callsign and track number), latitude/longitude position, height, ground speed, movement direction, aircraft type, runway and stand metadata (e.g., assigned runway and parking stand), and arrival/departure labels. These fields provide irregular positional samples (1–4 Hz), but they do not directly indicate which taxiway segment an aircraft occupies.

We therefore resample all trajectories to a uniform 1-second interval and apply a Hidden Markov Model (HMM) map-matching algorithm. The algorithm aligns aircraft positions to a directed taxiway network graph $G = (V, E)$, where nodes $v \in V$ represent intersections, runway entry points, and holding positions, and edges $e \in E$ represent taxiway segments with direction and length. Map matching is performed with a node-projection tolerance of $\delta_{\text{node}} = 5$ m and an edge-projection tolerance of $\delta_{\text{edge}} = 20$ m (consistent with taxiway widths of 15–25 m). The resulting network-aligned trajectories enable GIP detection based on shared segments, replacing coordinate-based proximity with topologically grounded interaction classification.

4.2. GIP Detection

Geometric Interaction Points identify locations where aircraft movements can constrain each other. We use two complementary detection modes to capture both structured taxiway interactions and unstructured apron encounters, but map all detected encounters onto one shared four-type vocabulary used throughout the paper: crossing (X), merge (M), head-on (H), and proximity (P).

Direct spatial GIPs arise from trajectory geometry. Two aircraft i and j form a spatial GIP candidate when their trajectories pass within a spatial proximity threshold within an overlapping time window:

$$d(i, j, t) = \|p_i(t) - p_j(t)\|_2 \le d_{\text{GIP}}, \quad \exists\, t \in T_i \cap T_j$$

where $p_i(t)$ is the position of aircraft i at time t, T_i is its operational time interval, and $d_{\text{GIP}} = 50$ m is the detection threshold. This threshold was set to match the taxiway separation standard while capturing encounters in unstructured apron areas where no shared-segment topology exists. Direct spatial GIPs are labeled as **proximity (P)** because they represent close encounters without a stable topological crossing or merge relation.

Topological GIPs arise from the network structure. We classify network-aligned interactions into three additional labels: **crossing (X)** for route intersections at the same node, **merge (M)** for aircraft converging onto the same downstream segment from different entries, and **head-on (H)** for opposing directions on the same segment. Co-directional following on the same segment is also mapped to **proximity (P)**, unifying it with apron-area near-encounters under a single analysis label for the motif view. For shared-segment following detection, a minimum segment length of $l_{\min} = 200$ m is required to exclude momentary spatial co-incidence at short connector segments.

4.3. SES Extraction

For each GIP involving two or more temporally proximate aircraft, we extract a **Spatio-temporal Encounter Session** capturing the behavioral signature of the interaction. An SES is defined over a temporal analysis window centered on the GIP passage time. For direct spatial GIPs, the window half-width is $w_{\text{spatial}} = 60$ s; for topological GIPs, the window extends $w_{\text{before}} = 30$ s before entry and $w_{\text{after}} = 30$ s after exit. A pre-filter discards candidate pairs whose passage times differ by more than $\Delta t_{\text{filter}} = 90$ s:

$$|t_i^{\text{pass}} - t_j^{\text{pass}}| \le \Delta t_{\text{filter}}, \quad \Delta t_{\text{filter}} = 90 \text{ s}$$

Within each analysis window, we extract behavioral descriptors: arrival order (first-to-pass vs. yielding), minimum separation $d_{\min}$, maximum closure velocity v_{closure}, speed reduction Δv, and total stop duration t_{stop}.

4.3.1. Interaction Severity Index

We compute an **Interaction Severity Index (ISI)** that integrates safety risk with operational efficiency loss into a single score for ranking and exploration across SES types:

$$\text{ISI} = w_r \cdot R_{\text{risk}} + w_e \cdot I_{\text{efficiency}}, \quad w_r + w_e = 1$$

The **risk exposure component** R_{risk} captures safety risk through separation distance and closure velocity:

$$R_{\text{risk}} = \alpha \cdot \text{norm}\left(\frac{1}{d_{\min} + \varepsilon}\right) + (1 - \alpha) \cdot \text{norm}(v_{\text{closure}})$$

where $d_{\min}$ is the minimum inter-aircraft separation distance, v_{closure} is the maximum relative closure velocity, $\varepsilon = 0.01$ prevents numerical instability when $d_{\min} \to 0$, and $\alpha = 0.5$. The term $1/(d_{\min} + \varepsilon)$ is evaluated after clamping $d_{\min}$ to $[d_{\text{safe}}, d_{\text{GIP}}]$, where d_{safe} is a conservative lower distance floor used only for bounded normalization and $d_{\text{GIP}} = 50$ m is the GIP detection radius, preventing extreme amplification of near-zero separations.

The **efficiency loss component** $I_{\text{efficiency}}$ captures operational disruption through deceleration and stop duration:

$$I_{\text{efficiency}} = \beta \cdot \text{norm}(\Delta v) + (1 - \beta) \cdot \text{norm}(t_{\text{stop}})$$

where Δv is the speed reduction relative to the free-flow speed and t_{stop} is the total duration of complete stops during the encounter, with $\beta = 0.5$.

We use equal weighting $w_r = w_e = 0.5$ as the default, treating risk and efficiency as complementary dimensions of equal operational priority. This configuration was established through iterative review with our domain experts (E1–E3, Section 3.3), who emphasized that low-risk but high-delay encounters can be just as operationally important as brief high-risk ones. The weights remain adjustable to reflect airport-specific policy differences. All norm$(\cdot)$ operations apply global min-max scaling over the full dataset.

4.4. Cascade Construction

Cascades capture how a local interaction propagates through a chain of coupled events. We first build an SES coupling graph with three edge types that encode different evidential strengths: **shared aircraft** for explicit flight continuity, **resource overlap** for short-range queueing or same-resource dependence, and **spatially reachable only** for temporally feasible topological continuation without a stronger anchor. The system's default analysis set, used in Sections 6 and 7, retains the exported 104-cascade configuration aligned with the motif fingerprints, while Section 7 reports sensitivity results for stricter and looser coupling modes.

4.4.1. Event Graph Construction

An edge from SES s_i to SES s_j is established when the two events are temporally feasible and at least one coupling predicate holds:

$$\Delta t_{ij} = t_j^{\text{start}} - t_i^{\text{end}} \in (0, \Delta t_{\text{couple}}], \quad \phi_{\text{shared}} \vee \phi_{\text{resource}} \vee \phi_{\text{spatial}}$$

where t_j^{start} and t_i^{end} are the start and end times of the respective SES, and $\Delta t_{\text{couple}} = 900$ s is the coupling time window.

The three coupling predicates are defined as follows.

- **Shared-aircraft coupling** (ϕ_{shared}): the two SES share at least one callsign, $\exists\, a \in A_i \cap A_j$. This is treated as the strongest evidence and assigned confidence 1.0. Because the shared flight is directly observed in both SES, we require only temporal feasibility.
- **Resource-overlap coupling** (ϕ_{resource}): no callsign is shared, but the two SES map to the same network node, the same parent edge, or a short topological separation that is consistent with localized queue spillback. In implementation, these edges are assigned confidence 0.7.
- **Spatial-reachability coupling** (ϕ_{spatial}): no stronger cue is present, but the downstream GIP is still reachable on the taxiway graph within a conservative distance bound $d_{\text{reach}} = 2000$ m and the same temporal window. These edges are assigned confidence 0.3 and are treated as the weakest evidence.

To separate evidence strength from reporting choice, we export four coupling modes: `strict` keeps only shared-aircraft edges; `extended` allows mixed paths but requires at least one shared-aircraft or strong resource anchor; the **default** configuration is the main analysis set used by the interface; and `unified` relaxes the path filter to provide an upper-bound reference.

We report coupling-mode sensitivity and targeted structural sanity checks in Section 7 to contextualize these assumptions and clarify how coupling definitions affect cascade coverage.

4.4.2. Cascade Extraction

From the coupling graph, we extract **cascade instances** as temporally ordered simple paths. We first enumerate source-to-sink paths and then remove strict subpaths so that only the longest unique paths are retained for analysis. The underlying coupling structure may still form a DAG when one upstream interaction fans out into multiple downstream interactions. Because the visual interface is organized around chronological replay, AirCascades materializes each retained branch as a separate linear instance. Shared prefixes may therefore appear in more than one instance, and all cascade counts reported in this paper are **path-instance counts** rather than connected-component counts. Each instance is characterized by its SES sequence, total event count, aircraft count, temporal span, spatial extent, coupling composition, and aggregate ISI statistics.

4.5. Motif Clustering

To surface recurring patterns among daily cascade instances, we encode each instance as a compact feature vector and cluster structurally similar instances into **motifs**.

4.5.1. Feature Vector

To align motif discovery with the glyph design in Section 5, each cascade instance is first converted into a **canonical interaction signature**. We traverse the SES sequence in time order, collapse consecutive duplicate interaction types, and retain the first three signature positions. Each position is encoded by a one-hot vector over the four analysis labels $\{X, M, H, P\}$, and shorter signatures are zero-padded. We then append five scalar descriptors for size, timing, and severity; together, these terms form the 17-dimensional representation in Table 1.

Table 1
Feature dimensions used for cascade representation.

Dim.	Category	Feature	Definition
1	Signature	Slot-1 = X	First canonical interaction type is Crossing
2	Signature	Slot-1 = M	First canonical interaction type is Merge
3	Signature	Slot-1 = H	First canonical interaction type is Head-on
4	Signature	Slot-1 = P	First canonical interaction type is Proximity
5	Signature	Slot-2 = X	Second canonical interaction type is Crossing
6	Signature	Slot-2 = M	Second canonical interaction type is Merge
7	Signature	Slot-2 = H	Second canonical interaction type is Head-on
8	Signature	Slot-2 = P	Second canonical interaction type is Proximity
9	Signature	Slot-3 = X	Third canonical interaction type is Crossing
10	Signature	Slot-3 = M	Third canonical interaction type is Merge
11	Signature	Slot-3 = H	Third canonical interaction type is Head-on
12	Signature	Slot-3 = P	Third canonical interaction type is Proximity
13	Summary	Event count	Number of SES events in the cascade instance
14	Summary	Mean inter-event lag	Mean time gap between consecutive SES events (s)
15	Summary	Total duration	End-to-end cascade duration (s)
16	Summary	Mean ISI	Mean Interaction Severity Index across all SES
17	Summary	Maximum ISI	Maximum ISI observed in any single SES

Before clustering, all features are standardized using z-score normalization:

$$\tilde{x}_k = \frac{x_k - \mu_k}{\sigma_k}$$

where μ_k and σ_k are the mean and standard deviation of feature k computed over the full cascade collection, ensuring that signature slots, temporal lags, and severity scores contribute comparably to the distance metric.

4.5.2. Cluster Selection and Configuration

We choose the number of clusters K using two criteria: the within-cluster sum of squares (WCSS) elbow and the mean silhouette coefficient:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i),\, b(i)\}}, \quad \bar{s} = \frac{1}{N}\sum_{i=1}^{N} s(i)$$

where $a(i)$ is the mean intra-cluster distance for instance i and $b(i)$ is the mean nearest-cluster distance. We evaluate $K \in [2, 10]$; both the WCSS elbow and peak $\bar{s}$ converge at $K = 9$. Clustering uses K-Means with $n_{\text{init}} = 10$ and random seed 42.

We chose K-Means rather than agglomerative hierarchical clustering because the 17-dimensional feature space does not show a natural dendrogram cutpoint and K-Means scales better to larger cascade collections. For presentation, each cluster is assigned a motif label from its modal canonical signature (e.g., **H-X-M**), while glyph color and summary statistics show within-cluster severity variation. On the default 104-cascade motif-mining set, this configuration yields nine motif families. Section 7 reports the silhouette score and motif-coverage changes under alternative coupling modes.

5. System Design

Building on the data abstractions and analytical tasks introduced in Section 3, we present AirCascades, a visual analytics system that supports exploration from overview to mechanism tracing. Figure 3 shows the complete interface.

5.1. Design Rationale

Our design addresses a fundamental tension in cascade analysis: controllers need both a global overview to identify critical incidents and detailed drill-down capabilities to understand propagation mechanisms. Two principles guided the interface design.

Progressive Disclosure. Following Shneiderman's information seeking mantra (Shneiderman, 1996), the interface is structured around three abstraction levels that mirror the analytical tasks: Level-1 provides macro-level overview for identifying significant cascades (T1; realized as G1, G2, G5), Level-2 enables meso-level pattern comparison for understanding structural similarities (T2; realized as G3, G5), and Level-3 supports micro-level narrative reconstruction for mechanism tracing (T3; realized as G4, G5, G6). Users can fluidly transition between levels without losing context.

Separation of Concerns. We employ coordinated multiple views (Roberts, 2007) that separate distinct analytical dimensions. Spatial context is isolated in the map view, structural patterns are abstracted into map-free glyphs, and temporal dynamics are visualized through narrative flows. This separation prevents visual clutter while enabling focused analysis of each dimension.

Alternatives Considered. We considered alternatives for the main components in Figure 3. A 2x2 grid gave equal space to all views, but it compressed the elongated airport map and interrupted the triage-to-replay workflow. This led to the triptych layout in Figure 3: View A supports ranked screening, View B preserves map space, and Views

Figure 3: The AirCascades system interface showing the coordinated macro-to-micro layout. (A) The Cascade Leaderboard supports cascade triage with ranking, summary metrics, and embedded sparklines. (B) The Geospatial View provides spatial context, route highlighting, and local zoom for selected cascades. (C) The Pattern View summarizes recurring structural families with map-free glyphs and instance cards. (D) The Micro Replay View exposes event-level speed curves, shared-aircraft links, and replay controls for detailed tracing.

C-D support pattern and mechanism analysis. Early map prototypes drew cascade paths directly on taxiway segments, but experts found junctions cluttered and direction hard to follow. We therefore separated the gray network base from the cyan propagation overlay. For motif families (View C), we tried small map overlays, but geographically different cascades with similar structures still looked unrelated. This led to map-free motif glyphs, which trade spatial specificity for structure-first comparison. For micro-level replay (View D), we considered a pure node-link timeline, but it did not show aircraft speed changes or make direct shared-aircraft continuity salient enough. We therefore used speed-curve cards with layered temporal connectors and shared-aircraft arcs, keeping behavioral evidence and continuity evidence visually separate.

5.2. System Architecture

AirCascades employs a **triptych layout** (Figure 3) reflecting the natural analysis workflow: a ranked cascade list on the left (View A) for triage, a central map (View B) for spatial context, and a right panel split between the Pattern View (View C) and detailed replay (View D). The triptych maximizes horizontal space for the map while establishing a clear analytical flow: *discover* → *locate* → *compare* → *explain*.

All views share a unified semantic color scheme: blue-gray tones for static network structures, cyan for event flows and temporal connections, a green-yellow-red gradient for risk severity, and purple hues for aircraft-specific encodings. Selections and filters are coordinated across views.

5.3. Visual Encoding Summary

Table 2 summarizes the core visual encodings across the four views. Rather than exhaustively listing every UI mark, it highlights the encodings most relevant to the analytical tasks and cross-view coordination.

5.4. Level-1: Macro-level Overview

The macro-level overview addresses T1 (*identify significant cascades*; G1, G2, G5) through two complementary views that provide ranking and spatial context.

5.4.1. Cascade Leaderboard

The **Cascade Leaderboard** (View A in Figure 3) presents all detected cascades as a sortable, filterable table where visual encodings are embedded directly into cells to support rapid scanning. Each row represents one **cascade instance**. Event count is encoded as a proportional bar chart for immediate size comparison, while risk severity is shown through both a numeric ISI value and a color-coded sparkline (Tufte, 2001) revealing how severity evolves across the event sequence. Purple dots overlaid on the sparkline indicate events where the same aircraft participates in consecutive interactions, surfacing the strongest continuity anchor without drill-down. Duration is summarized as a compact badge (e.g., "14m", "2h3m"), allowing analysts to assess temporal scale alongside severity and size.

A summary row at the top of View A displays aggregate statistics such as total cascade count, total events, average duration, and mean ISI. Together, these encodings support quick triage by helping analysts distinguish long but

Table 2
Visual encoding inventory across the four AirCascades views.

View	Key Encodings	Analytical Role
A: Leaderboard	Event-count bar; ISI sparkline; shared-aircraft marker	Supports cascade triage by showing size, severity evolution, and direct continuity anchors
B: Geospatial Map	Propagation path; SES nodes / heatmap	Reveals cascade route, propagation direction, local severity, and event concentration
C: Pattern View	Motif glyph; instance card	Supports structure-first comparison through canonical signatures, instance risk, and spatial footprint
D: Micro Replay	Dual speed curves; shared-aircraft arc; temporal connector / starglyph	Supports mechanism tracing through role behavior, direct propagation links, and inter-event spacing

low-severity cascades from shorter instances that escalate sharply or contain direct shared-aircraft anchors.

5.4.2. Geospatial Map View

When a cascade is selected, the **Geospatial View** (View B in Figure 3) reveals its spatial footprint using a dual-layer design that separates physical structure from propagation dynamics. The taxiway network is rendered as a semi-transparent gray base layer that provides geographic reference without visual dominance, while the selected cascade is drawn as a cyan Bezier path with directional arrows and numbered event markers. This separation addresses a critical design challenge: paths drawn directly on the network produced visual spikes at junctions and obscured direction. The final design makes *where the cascade occurred* and *how it propagated* visually distinguishable.

Layer toggles in View B allow users to overlay SES distributions, risk heatmaps, key network nodes, and aggregate flow patterns. A temporal strip at the bottom shows event density over time and supports time-range filtering that propagates across the coordinated views.

5.5. Level-2: Meso-level Pattern View

The meso-level **Pattern View** (View C in Figure 3) addresses T2 (*compare structural patterns*; G3, G5) by presenting motif families as **map-free glyphs** (Borgo, Kehrer, Chung, Maguire, Laramee, Hauser, Ward and Chen, 2013). Geographic maps are often less effective for structure-first comparison: cascades with similar interaction signatures but occurring at different airport locations can appear visually dissimilar and are hard to juxtapose. We therefore abstract each motif family into a topological glyph that encodes propagation structure without geographic coordinates.

5.5.1. Motif Glyph Design

Each glyph represents a family of structurally similar cascades using a dual-track metaphor in which two horizontal lines represent the trajectories of interacting aircraft pairs. The motif title is derived from the modal three-slot canonical signature introduced in Section 4.5 (for example, **H-X-M**). Crossing events are shown as intersecting strokes, merges as converging lines, head-on interactions as opposed arrows, and proximity events as closely aligned lines that approach without crossing. The color of each interaction symbol maps to its local ISI value using the unified risk color scale, enabling users to identify which interaction types tend to be most severe within a family.

5.5.2. Instance Strip

Below each glyph, an instance strip displays thumbnail representations of individual cascades belonging to that family. Instances are **sorted by risk severity** (highest ISI first), with border colors encoding risk levels. This design directs attention to the most dangerous cases within each pattern family, supporting the analytical goal of identifying high-risk anomalies.

Clicking a motif card filters View A and View B to show only cascades matching that structural pattern, linking pattern-level insight with geographic and temporal context.

5.6. Level-3: Micro-level Narrative View

The micro-level **Micro Replay View** (View D in Figure 3) addresses T3 (*reconstruct propagation mechanisms*; G4, G5, G6) by presenting individual cascade instances as **narrative flows** that reveal aircraft interactions over time. This design draws inspiration from Storyline visualizations (Tanahashi and Ma, 2012), adapting the concept to expose direct shared-aircraft continuity when it exists while keeping cascade timing visible.

5.6.1. Speed Curve Cards

Each SES event is represented as a card containing dual speed curves, one for each aircraft involved. The curves use a consistent purple color scheme: deep purple for the first-to-pass aircraft and light purple for the yielding aircraft. Header bars display aircraft metadata (callsign, aircraft type, arrival/departure status), providing context for understanding interaction dynamics. A replay button (D5) animates aircraft movement on the map, with a synchronized cursor on the speed curves showing the current playback position.

Adjacent to the speed curve cards, a **starglyph** provides a compact spatial summary of the cascade. The starglyph encodes the geographic path of the cascade as a star-shaped diagram: each branch represents a segment between consecutive GIP locations, with branch length proportional to

distance and branch color encoding local ISI severity. This helps analysts grasp cascade extent and severity distribution without opening the map.

5.6.2. Three-Layer Connection Design

Between adjacent event cards, we employ a three-layer visual hierarchy to encode different relationship types without overloading a single mark. A faint sky-blue band forms the **spatial reference layer**, indicating the main taxiway route in the background. On top of it, cyan connectors form the **temporal connection layer**: solid thick lines indicate intervals under 2 minutes, dashed lines indicate 2–5 minute gaps, and dotted thin lines indicate looser intervals over 5 minutes. Finally, purple arcs form the **shared-aircraft layer**, connecting adjacent events that involve the same callsign. These arcs appear only for direct aircraft continuity and can reveal the shared aircraft's intermediate trajectory on the map.

This layered design emerged from the need to show route context, temporal spacing, and direct continuity simultaneously. Rendering shared-aircraft arcs at the top makes the strongest propagation anchors visually salient, while the temporal connectors preserve cascade continuity even when no direct actor recurs.

5.6.3. Coordinated Trajectory Display

Clicking a speed curve card highlights the corresponding aircraft trajectories on the map (View B), with deep purple for the first-to-pass aircraft and light purple for the yielding aircraft. Clicking a shared-aircraft arc shows both connected events plus the shared aircraft's intermediate path. If its role changed (e.g., from yielding to first-to-pass), the trajectory uses a gradient color transition to show the reversal and help explain delay propagation.

5.7. Interaction and Coordination

The views are tightly coordinated to support fluid analytical workflows across levels. Selecting a row in the **Cascade Leaderboard** highlights the corresponding path in **View B** and populates **View D** with the selected replay. Clicking a motif card in **View C** filters the leaderboard and map to matching cascades, while interacting with speed-curve cards or shared-aircraft arcs in **View D** reveals detailed trajectories in the map. Temporal filtering in the map strip propagates to all views. These coordinated interactions allow users to move from overview triage to pattern comparison and then to mechanism tracing without resetting context.

6. Usage Scenarios

We illustrate AirCascades' analytical capabilities through two usage scenarios using real operational data from Beijing Capital International Airport (PEK). The dataset comprises 24 hours of ADS-B surveillance trajectories from November 2, 2023. Table 3 summarizes the default analysis set used in the interface and in the scenarios below; alternative coupling modes are compared separately in Section 7.

Table 3
Dataset statistics at each abstraction level.

Abstraction Level	Count	Description
Raw Trajectories	1,247	Aircraft movements during 24-hour period
GIPs	89	Geometric interaction points across taxiway network
SES	342	Spatio-temporal encounter sessions
Cascades	104	Default cascade instances (path-based counting)
Motifs	9	Structural pattern families

The following scenarios are representative analysis vignettes observed during the expert walkthrough sessions summarized in Section 7. They illustrate how AirCascades supports a typical workflow from global screening to pattern-level comparison and instance-level replay.

6.1. Scenario 1: Discovering High-Risk Propagation Patterns

E1 began with the Geospatial Map View and enabled the Key Nodes and Flow Paths layers to screen the overall risk distribution. The Key Nodes layer (Figure 4(a)) highlighted elevated-risk locations around apron entry/exit points and busy intersections, providing a compact shortlist of areas for deeper inspection.

The Flow Paths layer (Figure 4(b)) provided structural context by summarizing frequent propagation routes. E1 contrasted a compact, loop-like high-severity region in the west with lower-severity tree-like structures elsewhere, then used this contrast to prioritize which area to analyze next.

With a shortlist of candidate regions, E1 turned to the Pattern View (Figure 4(c)) to compare motif families without geographic distraction. The glyphs and severity summaries made it straightforward to identify a compact family with consistently higher severity than others.

E1 then used the instance strip to focus on the highest-severity examples within that family. Clicking an instance filtered the Leaderboard and updated the map to highlight only the corresponding SES events, allowing E1 to connect the pattern-level signature to its spatial footprint and to the specific cascades that instantiate it.

This vignette illustrates a top-down workflow supported by coordinated views: global screening (map layers) to pattern comparison (motif glyphs) to instance-level inspection (filtered leaderboard and map). The key payoff was not simply identifying another hotspot, but recognizing that the western area repeatedly produced a compact high-severity propagation structure. With alert logs or surveillance replay, analysts would have to review the cases manually; in AirCascades, the recurring family became visible before E1 opened detailed instances.

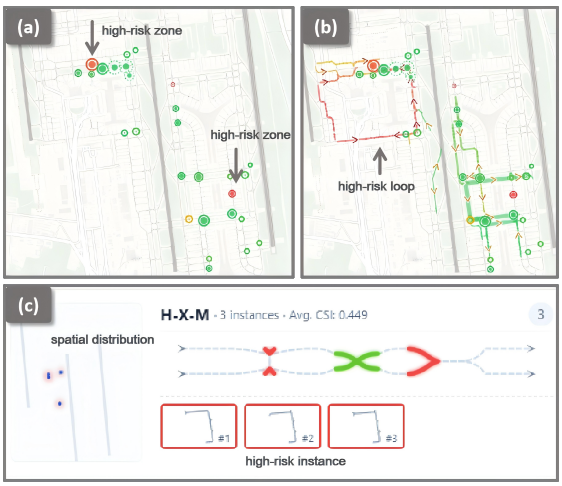

Figure 4: Macro-level pattern discovery. (a) The Key Nodes layer highlights elevated-risk locations concentrated around busy apron access points. (b) The Flow Paths layer reveals a compact high-severity propagation loop in the western area, contrasting with lower-severity structures elsewhere. (c) The Pattern View surfaces a compact high-risk motif family through map-free glyphs, spatial thumbnails, and highlighted instance cards.

6.2. Scenario 2: Tracing Propagation Through Shared Aircraft

E2 selected a Leaderboard cascade to trace the mechanism. The chosen instance belongs to the default analysis set and includes explicit shared-aircraft anchors, making it suitable for micro-level replay (Figure 5(a)).

In the Micro Replay view, E2 examined the speed curve cards arranged chronologically. Each card displayed dual speed profiles: the deep purple curve representing the first-to-pass aircraft and the light purple curve representing the yielding aircraft. The header bars showed aircraft metadata including callsign, aircraft type, and arrival/departure status.

E2 used the shared-aircraft arc to identify directly anchored adjacency between events. In this example, a purple arc connects the first two event cards and indicates that the same callsign (e.g., **CCA1119**) appears in both events. The arc, together with the role labels on the cards, supports quick identification of how one aircraft reappears as the carrier of continuity across multiple encounters.

To understand the spatial context, E2 examined the cascade path on the map (Figure 5(b)). The numbered circles marked GIP locations where interactions occurred, and the cyan arrows indicated propagation direction. The L-shaped path indicated that the cascade progressed from the central apron area toward the eastern runway.

E2 then clicked a speed curve card to inspect a single SES (Figure 5(c)). The map displayed both aircraft trajectories for that SES: the deep purple path labeled "first" for CCA1119 and the light purple path labeled "second" for the yielding arrival. The diamond marker indicated the GIP interaction point.

Finally, E2 clicked on the shared-aircraft arc to view both connected events simultaneously (Figure 5(d)). The map now displayed trajectories from both SES events, revealing how CCA1119's path intersected with two different arriving aircraft at different locations. This coordinated display enabled direct comparison of the shared aircraft's behavior across consecutive interactions.

This vignette illustrates a mechanism-oriented workflow: selecting a cascade, using shared-aircraft arcs to identify directly anchored links, and using coordinated map overlays to compare the shared aircraft's behavior across events. The key insight for E2 was that CCA1119 acted as the carrier of continuity across two encounters and changed its operational role across the cascade. Conventional surveillance replay can show both encounters, but reconstructing their connection requires manual callsign tracking and repeated temporal scrubbing; AirCascades surfaces that propagation carrier directly in a single coordinated view.

7. Evaluation

We evaluate AirCascades from three angles: coupling-mode sensitivity, structural sanity checks for high-priority cascades, and qualitative expert feedback. Unless otherwise stated, Sections 5–6 use the **default analysis set of 104 cascade instances**, which aligns with the exported motif fingerprints and provides adequate coverage for pattern-level comparison.

These three angles address complementary evaluation questions. Coupling sensitivity examines how strongly the reported cascade set depends on the edge-retention policy and whether the default configuration provides a defensible balance between coverage and evidential strength. The structural audit checks whether high-priority path-instances satisfy basic chronological and anchoring plausibility. The expert walkthrough investigates whether the coordinated views support effective post-hoc reasoning for domain practitioners. Together, these analyses assess pipeline validity, output quality, and practical utility.

7.1. Coupling Sensitivity

Table 4 compares four exported coupling modes. The `strict` mode keeps only explicit shared-aircraft edges; `extended` allows mixed paths but requires at least one shared-aircraft or strong resource anchor; the **default** configuration is the main analysis set used throughout the interface; and `unified` is a looser upper-bound reference that applies only a confidence threshold.

Coupling choice materially affects cascade coverage. The strongest-evidence subsets are extremely small (`strict`: 3, `extended`: 5) and serve mainly as high-precision references. We therefore use the default configuration (104) as the main analysis set because it preserves a substantially broader pattern inventory while remaining more conservative than the looser `unified` reference (113).

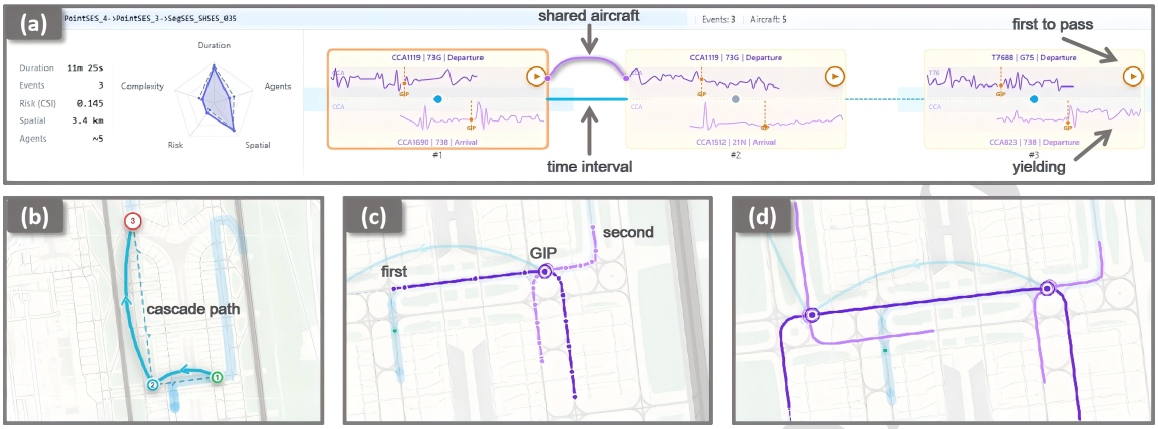

Figure 5: Micro-level propagation tracing. (a) The Micro Replay view shows speed-curve cards, shared-aircraft arcs, and temporal alignment cues across the selected cascade. (b) The corresponding map view displays the cascade route with numbered GIP markers. (c) Selecting a single SES card reveals the paired trajectories and encounter location for that event. (d) Selecting a shared-aircraft link reveals the connected events simultaneously for cross-event comparison.

Table 4
Cascade counts and evidence composition under four coupling modes.

Mode	Count	Shared	Resource	Spatial
Strict	3	1.000	0.000	0.000
Extended	5	0.767	0.233	0.000
Default	104	0.150	0.157	0.693
Unified	113	0.310	0.081	0.609

Evidence composition also shifts across modes. In the default set, 38.46% of cascades contain at least one shared-aircraft anchor, while 38.46% are spatial-only paths. In the unified set, shared-anchor coverage rises to 64.60% and spatial-only paths fall to 27.43%. We therefore report these modes as a sensitivity band rather than interchangeable estimates of a single "true" cascade count. The default mode is best interpreted as an analysis construct that retains weaker but operationally plausible continuities, whereas the strict and extended modes serve as higher-confidence lower-coverage references.

Motif mining is performed on the default 104-cascade set, yielding 9 motif families with a silhouette score of 0.254. This indicates only moderate separation: suitable for exploratory pattern comparison, but insufficient for strong claims about fixed cluster boundaries.

7.2. Structural Audit

We conducted an author-side structural audit of the top 25 default cascades ranked by average ISI to check whether the path-instances emphasized by the interface satisfy basic plausibility criteria. This audit is a targeted sanity check rather than a gold-standard causality validation. It focuses on three simple properties: temporal monotonicity of the event sequence, the presence of at least one non-spatial anchor (shared-aircraft or resource-overlap evidence), and explicit shared-aircraft overlap detectable from callsign intersections in adjacent SES pairs.

All 25 audited cascades are temporally monotone, indicating that the path extraction step preserves basic chronological consistency. More importantly, 16 of the 25 audited paths contain at least one shared-aircraft or resource-overlap anchor, and 10 contain at least one explicit shared-aircraft overlap in adjacent SES pairs. These anchored cases are the ones most directly aligned with the propagation semantics emphasized throughout the paper.

The remaining 9 of the 25 audited cascades are spatial-only under the default configuration. We do not interpret them as confirmed causal chains; rather, they represent weaker but operationally plausible propagation hypotheses that become available only when the analysis favors coverage over strict continuity evidence. We additionally checked the `strict` subset as a high-precision reference, and all 3 strict cascades have explicit shared callsign overlap on all adjacent pairs (3/3). Taken together, these results clarify the trade-off surfaced in Section 7.1: strong shared-aircraft evidence is precise but very sparse, whereas the default set broadens structural coverage with greater uncertainty for spatial-only paths.

7.3. Expert Walkthrough Feedback

We conducted walkthrough-based interviews with the same three experts described in Section 3.3. Using the same 24-hour PEK dataset analyzed in Section 6, each session asked the experts to inspect high-risk regions, compare motif families, and trace a selected cascade through the micro replay before participating in a semi-structured interview lasting about 45 minutes. Because these participants also helped shape the abstractions and visual encodings, we report the results as **qualitative expert feedback** rather than an independent summative user study.

Table 5
Structural audit summary for the top 25 default cascades ranked by average ISI.

Audit check	Result	Interpretation
Top-ranked default cascades reviewed	25	High-priority subset used for targeted plausibility checking
Temporally monotone event sequences	25/25	No audited path violated chronological ordering
Cascades with at least one shared-aircraft or resource anchor	16/25	Most audited paths contain evidence stronger than pure spatial reachability
Cascades with at least one explicit shared-aircraft overlap	10/25	A substantial subset contains directly observed continuity
Adjacent SES pairs with explicit shared-aircraft overlap	10/68	Direct continuity is informative but not pervasive in the default mode
Strict-reference cascades with explicit shared overlap on all adjacent pairs	3/3	Shared-only paths provide a high-precision reference

At the macro level, the experts reported that the progressive-disclosure workflow supports post-hoc reasoning from hotspot overview to cascade replay. They found the **Cascade Leaderboard** useful for triage because the combination of event count, ISI ranking, and embedded sparklines made it easy to distinguish large but routine cascades from shorter cases with sharp severity escalation. They also highlighted the layer toggles in the **Geospatial View** as a practical mechanism for separating spatial context, event density, and propagation paths without losing the selected cascade.

At the meso level, the experts considered the **map-free motif glyphs** effective for structure-first comparison because they made repeated propagation forms visible without requiring geographic alignment. This was especially valuable when comparing cascades that occurred in different parts of the surface but shared the same interaction progression. The main cost was a short learning curve for the X/M/H/P notation: all three experts reported that the symbols became readable after orientation, but none viewed them as immediately self-explanatory on first exposure.

At the micro level, the **shared-aircraft arcs** were consistently mentioned as the most informative encoding for identifying directly anchored links between adjacent events. The experts also found the role labels and speed curves helpful for interpreting who imposed delay, who yielded, and how those roles shifted across the cascade. Several comments suggested that the value of the micro replay lies less in replaying motion alone than in making continuity explicit: shared-aircraft anchors were trusted most, while purely temporal or topological continuity cues were seen as useful but more interpretive.

The walkthroughs also surfaced practical constraints rather than uniformly positive reactions. First, the experts questioned whether the fixed motif count ($K = 9$) would remain suitable under heavier traffic or at airports with different topologies. Second, they were more confident in cascades containing explicit shared-aircraft anchors than in spatial-only paths, reinforcing the need to interpret the default cascade set as an analytical aid rather than as ground-truth causality. Third, they emphasized that the current implementation is best suited for **post-hoc analysis** such as training, debriefing, and safety review; using it for real-time monitoring would require additional work on latency control, alert management, and threshold calibration. Overall, the walkthroughs suggest that AirCascades is most valuable when analysts move from broad situational screening to mechanism tracing within one coordinated workspace.

8. Discussion

8.1. Design Reflections

Our iterative design process yielded several insights applicable to visualizing propagation phenomena in constrained spatial networks.

Separating network structure from propagation dynamics proved essential. Early prototypes that rendered cascade paths directly on the taxiway network produced visual clutter at junctions. The dual-layer approach uses a subdued network base and a prominent propagation overlay, allowing users to perceive both spatial context and temporal flow without interference.

The three-layer connection design in the Micro Replay view emerged from the need to encode multiple relationship types simultaneously. By assigning distinct visual layers to spatial reference (bottom), temporal intervals (middle), and shared-aircraft links (top), we avoid overloading any single channel while maintaining clear semantic separation.

Map-free abstraction for pattern comparison addresses a limitation of geographic visualization: structurally similar cascades occurring at different locations are hard to compare on a map. The motif glyph design sacrifices spatial specificity to enable structure-first comparison, a trade-off that experts found valuable for identifying recurring patterns.

8.2. Implications

For the **visual analytics community**, AirCascades shows how cascade analysis can benefit from combining domain-constrained event extraction with representation choices tailored to a replay-oriented workflow. In particular, the path-instance formulation, canonical signature abstraction, and coordinated macro-to-micro views offer a template for studying propagation in settings where influence may be carried by explicit actor continuity and topology-aware couplings rather than aggregate flow alone.

For the **airport operations community**, the system reframes post-hoc surface analysis around propagation chains

instead of isolated alerts. This shift may support training, debriefing, and procedure review by making it easier to connect local interactions with downstream consequences, even when those consequences unfold across different parts of the surface network.

8.3. Limitations

Several limitations should be acknowledged. First, our evaluation uses data from a **single airport and day** together with a **small expert sample** (n=3). Although the three experts contribute a combined 26 years of experience and represent complementary operational and system perspectives, broader validation across airports, traffic regimes, and time spans is still needed before claiming strong generalizability.

Second, the walkthrough participants also contributed to the formative design process that shaped the abstractions, interaction rules, and visual encodings. Their feedback is therefore highly informative for assessing domain fit, but it also introduces a **confirmation-bias risk** compared with an evaluation using independent analysts who had no prior involvement in the system design.

Third, the reported cascade set remains dependent on the **coupling definition** and on the analytical abstractions built on top of it. As Section 7 shows, strict shared-aircraft-only coupling yields only 3 cascades, whereas looser configurations yield 104–113 path-instances; similarly, the 9-motif partition offers a useful exploratory summary but only moderate cluster separation. The reported cascades and motifs should therefore be interpreted as analysis constructs that trade evidential strength for broader structural coverage, not as ground-truth causal chains or uniquely correct clusters.

8.4. Transferability to Other Airports

Applying AirCascades to another airport would require local calibration, but not a change to the overall analysis workflow. The required adjustments fall into three groups. First, at the **data collection** level, the target airport should provide ADS-B trajectories and a taxiway network with local node, edge, runway, and stand definitions. Second, at the **data processing** level, map-matching tolerances, GIP detection thresholds, SES temporal windows, cascade coupling settings, ISI ranges, and motif-clustering parameters should be checked against local taxiway widths, traffic density, controller procedures, and safety priorities. Third, at the **user interface** level, map aspect ratio, zoom levels, and glyph density may need adjustment to fit the geometry and traffic volume of the target airport. These are calibration steps, not changes to the core SES–cascade–motif abstraction. We therefore view AirCascades as transferable, with broader multi-airport validation left for future work.

8.5. Future Directions

Future work should test multi-day, multi-airport collections, calibrate coupling and ISI parameters against controller-annotated reference cases, and support analyst-in-the-loop refinement of cascade and motif definitions. These extensions would strengthen transferability while preserving the paper's focus on post-hoc propagation analysis.

9. Conclusion

We presented AirCascades, a visual analytics system for understanding cascading taxiway interactions at airports. Our approach bridges raw surveillance trajectories and operationally meaningful propagation analysis through a semantics-driven data abstraction pipeline that identifies interaction events, links them into cascade instances through evidence-weighted operational couplings, and groups these instances into structural motif families.

The system's multi-level visual design supports progressive exploration: a macro-level overview for identifying significant cascades and high-risk network regions, a meso-level pattern view for comparing structural signatures across cascade families, and a micro-level narrative view for reconstructing propagation mechanisms through synchronized speed curves and shared-aircraft connection arcs.

Coupling-sensitivity analysis shows that cascade coverage changes materially under different coupling definitions, while a focused structural audit and qualitative expert feedback suggest that AirCascades helps analysts inspect propagation dynamics that can be difficult to perceive from raw surveillance data alone. The shared-aircraft visualization explicitly surfaces the strongest directly observed links, while the map-free motif glyphs enable structure-first comparison that geographic displays do not readily support.

Future work includes multi-day deployment across airports with different topologies, coupling-parameter calibration against controller-annotated reference cases, contextual factors such as weather and schedule disruptions, and analyst-in-the-loop refinement of cascade interpretations.

CRediT authorship contribution statement

Guoqiang Wang: Writing - original draft, Writing - review & editing, Software, Visualization, Investigation, Formal analysis, Data curation, Methodology, Conceptualization. **Xiaolin Wen:** Writing - review & editing, Supervision, Methodology, Conceptualization. **Fengjie Wang:** Writing - review & editing, Methodology, Conceptualization. **Luge Yang:** Writing - review & editing, Validation. **Qipeng Wang:** Writing - review & editing, Validation. **Min Zhu:** Writing - review & editing, Project administration, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the National Key R&D Program of China (No. 2023YFB2604100). The authors thank the participating domain experts for their time, operational insights, and feedback during the formative study and walkthroughs.

References

Al-Dohuki, S., Kamw, F., Zhao, Y., Yang, J., Li, X., Wu, Y., Ye, X., Chen, W., Ma, C., Wang, F., 2017. SemanticTraj: A new approach to interacting with massive taxi trajectories. IEEE Transactions on Visualization and Computer Graphics 23, 11–20. doi:10.1109/TVCG.2016.2598416.

Albrecht, G.H., Lee, H.T., Pang, A., 2012. Visual analysis of air traffic data using aircraft density and conflict probability, in: Infotech@Aerospace 2012. doi:10.2514/6.2012-2540.

Ali, H., Delair, R., Pham, D.T., Alam, S., Schultz, M., 2020. Dynamic hot spot prediction by learning spatial-temporal utilization of taxiway intersections, in: 2020 International Conference on Artificial Intelligence and Data Analytics for Air Transportation (AIDA-AT), pp. 1–10. doi:10.1109/AIDA-AT48540.2020.9049186.

Andrienko, G., Andrienko, N., Bak, P., Keim, D.A., Wrobel, S., 2013. Visual Analytics of Movement. Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-37583-5.

Andrienko, G., Andrienko, N., Chen, W., Maciejewski, R., Zhao, Y., 2017. Visual analytics of mobility and transportation: State of the art and further research directions. IEEE Transactions on Intelligent Transportation Systems 18, 2232–2249. doi:10.1109/TITS.2017.2683539.

Andrienko, N., Andrienko, G., 2013. Visual analytics of movement: An overview of methods, tools and procedures. Information Visualization 12, 3–24. doi:10.1177/1473871612457601.

Bach, B., Dragicevic, P., Archambault, D., Hurter, C., Carpendale, S., 2017. A descriptive framework for temporal data visualizations based on generalized space-time cubes. Computer Graphics Forum 36, 36–61. doi:10.1111/cgf.12804.

Borgo, R., Kehrer, J., Chung, D.H.S., Maguire, E., Laramee, R.S., Hauser, H., Ward, M., Chen, M., 2013. Glyph-based visualization: Foundations, design guidelines, techniques and applications, in: Eurographics 2013: State of the Art Reports, pp. 39–63. doi:10.2312/conf/EG2013/stars/039-063.

Chen, W., Guo, F., Wang, F.Y., 2015. A survey of traffic data visualization. IEEE Transactions on Intelligent Transportation Systems 16, 2970–2984. doi:10.1109/TITS.2015.2436897.

Deng, Z., Weng, D., Chen, J., Liu, R., Wang, Z., Bao, J., Zheng, Y., Wu, Y., 2020. AirVis: Visual analytics of air pollution propagation. IEEE Transactions on Visualization and Computer Graphics 26, 800–810. doi:10.1109/TVCG.2019.2934670.

Deng, Z., Weng, D., Liang, Y., Bao, J., Zheng, Y., Schreck, T., Xu, M., Wu, Y., 2022. Visual cascade analytics of large-scale spatiotemporal data. IEEE Transactions on Visualization and Computer Graphics 28, 2486–2499. doi:10.1109/TVCG.2021.3071387.

Deng, Z., Weng, D., Liu, S., Tian, Y., Xu, M., Wu, Y., 2023. A survey of urban visual analytics: Advances and future directions. Computational Visual Media 9, 3–39. doi:10.1007/s41095-022-0275-7.

Feng, Z., Zhu, F., Wang, H., Hao, J., Yang, S.H., Zeng, W., Qu, H., 2024. HoLens: A visual analytics design for higher-order movement modeling and visualization. Computational Visual Media 10, 1079–1100. doi:10.1007/s41095-023-0392-y.

Ferreira, A., Afonso, A.P., Ferreira, L., Vaz, R., 2022. Visual analytics of trajectories with RoseTrajVis. Big Data Research 27, 100294. doi:10.1016/j.bdr.2021.100294.

Ferreira, N., Poco, J., Vo, H.T., Freire, J., Silva, C.T., 2013. Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. IEEE Transactions on Visualization and Computer Graphics 19, 2149–2158. doi:10.1109/TVCG.2013.226.

Fleurquin, P., Ramasco, J.J., Eguiluz, V.M., 2013. Systemic delay propagation in the US airport network. Scientific Reports 3, 1159. doi:10.1038/srep01159.

Guo, H., Wang, Z., Yu, B., Zhao, H., Yuan, X., 2011. TripVista: Triple perspective visual trajectory analytics and its application on microscopic traffic data at a road intersection, in: 2011 IEEE Pacific Visualization Symposium, pp. 163–170. doi:10.1109/PACIFICVIS.2011.5742386.

Hu, D., Chen, L., Fang, H., Fang, Z., Li, T., Gao, Y., 2024. Spatio-temporal trajectory similarity measures: A comprehensive survey and quantitative study. IEEE Transactions on Knowledge and Data Engineering 36, 2191–2212. doi:10.1109/TKDE.2023.3323535.

Hurter, C., Ersoy, O., Telea, A., 2013. Smooth bundling of large streaming and sequence graphs, in: 2013 IEEE Pacific Visualization Symposium, pp. 41–48. doi:10.1109/PacificVis.2013.6596126.

ICAO, 2004. Manual of Advanced Surface Movement Guidance and Control Systems (A-SMGCS). Technical Report Doc 9830. International Civil Aviation Organization.

ICAO, 2016. Doc 4444: Procedures for Air Navigation Services — Air Traffic Management. International Civil Aviation Organization.

Jalilvand, A., Milhomem, L., Paulovich, F.V., 2025. SeRViz: A visual analytics system for the analysis of sequential rules and its application to airport ground handling operations. Journal of Visualization 28, 1169–1189. doi:10.1007/s12650-025-01088-z.

Jiang, Y., Xue, Q., Wang, Y., Cai, M., Zhang, H., Li, Y., 2021. Traffic congestion mechanism in mega-airport surface: A cellular automaton approach. Physica A: Statistical Mechanics and its Applications 577, 126080. doi:10.1016/j.physa.2021.126080.

Kuchar, J.K., Yang, L.C., 2000. A review of conflict detection and resolution modeling methods. IEEE Transactions on Intelligent Transportation Systems 1, 179–189. doi:10.1109/6979.898217.

von Landesberger, T., Bremm, S., Andrienko, N., Andrienko, G., Tekušová, M., 2012. Visual analytics methods for categoric spatio-temporal data, in: 2012 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 183–192. doi:10.1109/VAST.2012.6400553.

von Landesberger, T., Brodkorb, F., Roskosch, P., Andrienko, N., Andrienko, G., Kerren, A., 2016. MobilityGraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering. IEEE Transactions on Visualization and Computer Graphics 22, 11–20. doi:10.1109/TVCG.2015.2468111.

von Landesberger, T., Kuijper, A., Schreck, T., Kohlhammer, J., van Wijk, J.J., Fekete, J.D., Felkel, D.W., 2011. Visual analysis of large graphs: State-of-the-art and future research challenges. Computer Graphics Forum 30, 1719–1749. doi:10.1111/j.1467-8659.2011.01898.x.

Li, J., Chen, S., Zhang, K., Andrienko, G., Andrienko, N., 2019. COPE: Interactive exploration of co-occurrence patterns in spatial time series. IEEE Transactions on Visualization and Computer Graphics 25, 2554–2567. doi:10.1109/TVCG.2018.2851227.

Lu, M., Lai, C., Ye, T., Liang, J., Yuan, X., 2017. Visual analysis of multiple route choices based on general GPS trajectories. IEEE Transactions on Big Data 3, 234–247. doi:10.1109/TBDATA.2017.2667700.

Ma, J., Zhou, J., Liang, M., Delahaye, D., 2022. Data-driven trajectory-based analysis and optimization of airport surface movement. Transportation Research Part C: Emerging Technologies 145, 103902. doi:10.1016/j.trc.2022.103902.

Maciejewski, R., Rudolph, S., Hafen, R., Abusalah, A., Yakout, M., Ouzzani, M., Cleveland, W.S., Grannis, S.J., Ebert, D.S., 2010. A visual analytics approach to understanding spatiotemporal hotspots. IEEE Transactions on Visualization and Computer Graphics 16, 205–220. doi:10.1109/TVCG.2009.100.

Monroe, M., Lan, R., Lee, H., Plaisant, C., Shneiderman, B., 2013. Temporal event sequence simplification. IEEE Transactions on Visualization and Computer Graphics 19, 2227–2236. doi:10.1109/TVCG.2013.200.

Pi, M., Yeon, H., Son, H., Jang, Y., 2021. Visual cause analytics for traffic congestion. IEEE Transactions on Visualization and Computer Graphics 27, 2186–2201. doi:10.1109/TVCG.2019.2940580.

Roberts, J.C., 2007. State of the art: Coordinated and multiple views in exploratory visualization, in: Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007), pp. 61–71. doi:10.1109/CMV.2007.20.

Scheepens, R., Willems, N., van de Wetering, H., Andrienko, G., Andrienko, N., van Wijk, J.J., 2011. Composite density maps for multivariate trajectories. IEEE Transactions on Visualization and Computer Graphics 17, 2518–2527. doi:10.1109/TVCG.2011.181.

Sedlmair, M., Meyer, M., Munzner, T., 2012. Design study methodology: Reflections from the trenches and the stacks. IEEE Transactions on Visualization and Computer Graphics 18, 2431–2440. doi:10.1109/TVCG.2012.213.

Shneiderman, B., 1996. The eyes have it: A task by data type taxonomy for information visualizations, in: Proceedings of the IEEE Symposium on

Visual Languages, pp. 336–343. doi:10.1109/VL.1996.545307.

Tanahashi, Y., Ma, K.L., 2012. Design considerations for optimizing storyline visualizations. IEEE Transactions on Visualization and Computer Graphics 18, 2679–2688. doi:10.1109/TVCG.2012.212.

Tian, W., Zhou, X., Yin, J., Li, Y., Zhang, Y., 2024. Identification of key risk hotspots in mega-airport surface based on Monte Carlo simulation. Aerospace 11, 254. doi:10.3390/aerospace11040254.

Tufte, E.R., 2001. The Visual Display of Quantitative Information. 2 ed., Graphics Press.

Vaddi, S., Kwan, J., Fong, A., Cheng, V.H.L., 2012. Deterministic and probabilistic conflict detection algorithms for NextGen airport surface operations, in: AIAA Guidance, Navigation, and Control Conference. doi:10.2514/6.2012-4974.

Wang, F., Chen, W., Wu, F., Zhao, Y., Hong, H., Gu, T., Wang, L., Liang, R., Bao, H., 2014. A visual reasoning approach for data-driven transport assessment on urban roads, in: 2014 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 103–112. doi:10.1109/VAST.2014.7042483.

Wang, Z., Lu, M., Yuan, X., Zhang, J., van de Wetering, H., 2013. Visual traffic jam analysis based on trajectory data. IEEE Transactions on Visualization and Computer Graphics 19, 2159–2168. doi:10.1109/TVCG.2013.228.

Wu, Y., Weng, D., Deng, Z., Bao, J., Xu, M., Wang, Z., Zheng, Y., Ding, Z., Chen, W., 2021. Towards better detection and analysis of massive spatiotemporal co-occurrence patterns. IEEE Transactions on Intelligent Transportation Systems 22, 3387–3402. doi:10.1109/TITS.2020.2983226.

Xia, Z., Zheng, B., Wan, J., Zhu, X., 2019. Recognition algorithm and risk assessment of airport hotspots. Journal of Shanghai Jiao Tong University (Science) 24, 769–774. doi:10.1007/s12204-019-2110-6.

Yin, J., He, Y., Ma, Y., Qiao, P., Liu, X., 2025. Data-driven recognition of hot taxiway segments for assessing airport surface traffic efficiency. Journal of Air Transport Management 127, 102822. doi:10.1016/j.jairtraman.2025.102822.

Zeng, W., Fu, C.W., Müller Arisona, S., Erath, A., Qu, H., 2014. Visualizing mobility of public transportation system. IEEE Transactions on Visualization and Computer Graphics 20, 1833–1842. doi:10.1109/TVCG.2014.2346893.

Zeng, W., Fu, C.W., Müller Arisona, S., Erath, A., Qu, H., 2016. Visualizing waypoints-constrained origin-destination patterns for massive transportation data. Computer Graphics Forum 35, 95–107. doi:10.1111/cgf.12778.

Zhang, X., Zhong, S., Mahadevan, S., 2022. Airport surface movement prediction and safety assessment with spatial-temporal graph convolutional neural network. Transportation Research Part C: Emerging Technologies 144, 103873. doi:10.1016/j.trc.2022.103873.

Zhou, H., Lai, P., Sun, Z., Chen, X., Chen, Y., Wu, H., Wang, Y., 2025. AdaMotif: Graph simplification via adaptive motif design. IEEE Transactions on Visualization and Computer Graphics 31, 688–698. doi:10.1109/TVCG.2024.3456321.

Zhuge, J., Liang, H., Zhang, Y., Li, S., Yang, X., Wu, J., 2024. Aircraft ground taxiing deduction and conflict early warning method based on control command information. Transportation Research Record 2678, 684–704. doi:10.1177/03611981241242360.

Zohrevandi, E., Vrotsou, K., Westin, C.A.L., Lundberg, J., Ynnerman, A., 2024. Design of a real-time visual analytics decision support interface to manage air traffic complexity, in: IEEE VIS, pp. 301–305. doi:10.1109/VIS55277.2024.00068.

Declaration of Competing Interests

Manuscript Number: VISINF-D-26-00074

Manuscript Title: AirCascades: A Visual Analytics System for Cascading Taxiway Interaction Patterns

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Guoqiang Wang, Xiaolin Wen, Fengjie Wang, Luge Yang, Qipeng Wang, Min Zhu