

Contrail-to-Flight Attribution Using Ground Visible Cameras and Flight Surveillance Data

Ramon Dalmau, Gabriel Jarry & Philippe Very
Aviation Sustainability Unit (ASU)
EUROCONTROL
Brétigny-Sur-Orge, France

Abstract—Aviation’s non-CO₂ effects, particularly contrails, are a significant contributor to its climate impact. Persistent contrails can evolve into cirrus-like clouds that trap outgoing infrared radiation, with radiative forcing potentially comparable to or exceeding that of aviation’s CO₂ emissions. While physical models simulate contrail formation, evolution and dissipation, validating and calibrating these models requires linking observed contrails to the flights that generated them, a process known as contrail-to-flight attribution. Satellite-based attribution is challenging due to limited spatial and temporal resolution, as contrails often drift and deform before detection. In this paper, we evaluate an alternative approach using ground-based cameras, which capture contrails shortly after formation at high spatial and temporal resolution, when they remain thin, linear, and visually distinct. Leveraging the ground visible camera contrail sequences (GVCCS) dataset, we introduce a modular framework for attributing contrails observed using ground-based cameras to forecasted contrails derived from aircraft surveillance and meteorological data. The framework accommodates multiple geometric representations and distance metrics, incorporates temporal smoothing, and enables flexible probability-based assignment strategies. This work establishes a strong baseline and provides a modular framework for future research in linking contrails to their source flight.

Keywords—contrails; environmental impact; aviation; observations & forecast.

I. INTRODUCTION

Aviation contributes to climate change not only through CO₂ emissions but also through non-CO₂ effects, including nitrogen oxides (NO_x), water vapour, aerosols, and contrails, ice-crystal clouds formed by aircraft at cruising altitudes. Persistent contrails can spread into extensive cirrus-like formations, trapping outgoing long-wave radiation and warming the planet. Recent studies suggest that the radiative forcing of contrail cirrus may be comparable to that of aviation CO₂ emissions [1], [2], though the exact magnitude depends on the metric used [3].

Quantifying the climate impact of contrails, however, remains a major challenge. Their lifecycle depends on interconnected processes: ice nucleation, crystal growth, wind-driven dispersion, and interactions with natural clouds, all highly sensitive to atmospheric conditions. Small variations in temperature and relative humidity with respect to ice determine whether a contrail dissipates or persists. The radiative impact of contrails also depends on their optical properties, such as the size and number of ice crystals, which influence how much sunlight they reflect or infrared radiation they trap. Combined with day–night differences in radiative forcing (cooling from

sunlight during the day versus warming from infrared trapping at night), contrails’ overall climate effect varies widely depending on these conditions, making it difficult to model.

Most studies of contrail impacts rely on models such as CoCiP [4], which are often described as plume models or Lagrangian models. These simulate contrail formation based on aircraft parameters and atmospheric conditions, but they depend heavily on parametrisations of key physical mechanisms, including ice nucleation, crystal growth, contrail evolution, and optical properties. As a result, their accuracy is limited by the simplifications inherent in representing these microphysical processes, particularly the role of relative humidity [5]. Observations from satellites [6] and ground-based imagery [7] provide a crucial way to test and improve these models. However, this requires one essential step: linking each observed contrail to the specific flight that produced it. Without such attribution, it is impossible to assess whether a model’s predictions about contrail occurrence, persistence, or optical properties match reality. In short, contrail-to-flight attribution transforms raw imagery into actionable evidence for model calibration and validation.

Satellite imagery, however, has inherent limitations. By the time contrails are visible at satellite resolution, they have often drifted, deformed, or merged with other clouds, making alignment with flight trajectories highly uncertain. Ground-based cameras offer a promising alternative. They capture contrails just moments after formation, at high spatial and temporal resolution, when they are still thin, linear, and visually distinct. When paired with Automatic Dependent Surveillance–Broadcast (ADS-B) and meteorological data to simulate forecasted contrails, ground-based imagery opens a practical path toward precise contrail-to-flight attribution.

In this paper, we introduce a modular framework for contrail-to-flight attribution using ground-based visible cameras and evaluate it on the Ground Visible Camera Contrail Sequences (GVCCS) dataset [8], [9]. GVCCS contains video sequences recorded by a ground-based camera in Brétigny-sur-Orge, France, with detailed annotations for each contrail. Human labellers drew polygons directly on the images from the ground-based camera to capture contrail shapes, and each contrail was assigned a unique identifier that remains consistent across frames, enabling accurate tracking over time. Whenever possible, contrails were also linked to the flights that produced them, making the dataset ideally suited for testing



and refining contrail-to-flight attribution algorithms.

The rest of this paper is structured as follows. Section II reviews related work on contrail-to-flight attribution. Section III describes the data used in our study. Section IV presents the components of our modular framework, and Section V details the experimental setup. Section VI reports the results, and Section VII concludes with key findings and future directions.

II. STATE OF THE ART

Physical models of contrails, such as CoCiP [4] and APCEMM [10], provide detailed predictions of contrail formation and evolution, but their accuracy remains uncertain due to limited observational validation [5]. Observational methods, in turn, detect contrails in satellite [6] or ground-based [7] imagery, but without a reliable way to attribute contrails to their source flights, these observations cannot be linked to flight-specific parameters and used for systematic model-observation comparison. Contrail-to-flight attribution is therefore a critical step for closing the gap between physics-based and observational approaches in contrail science.

Satellite imagery has long been the primary source of observational data, but it has inherent limitations when used for contrail-to-flight attribution. Geostationary satellites offer continuous temporal coverage (5–15 min refresh) but at relatively coarse spatial resolution (0.5–2 km), meaning contrails are only visible once they have grown thick and displaced under wind [11]. Polar-orbiting satellites such as MODIS provide sharper images but only infrequent snapshots. In both cases, attribution becomes ambiguous as contrails drift, deform, or merge with surrounding clouds [12]. Furthermore, to the best of our knowledge, no human-labelled dataset exists that links contrails in satellite images to their source flights at scale, making direct evaluation impractical.

Synthetic datasets have been created to overcome this lack of ground truth. For instance, [6] generated contrails using CoCiP and overlaid them on GOES-16 imagery, enabling the first systematic contrail instance segmentation experiments. [13] extended this idea with SynthOpenContrails, which provides synthetic contrails linked to flights, creating the first benchmark for contrail-to-flight attribution algorithm comparison. While synthetic datasets are invaluable for algorithm development, their algorithmically generated ground truth prevents them from serving as a definitive benchmark for evaluation.

Ground-based cameras offer a complementary observational perspective. They capture contrails shortly after formation, when they remain thin, linear, and visually distinct. This high temporal and spatial resolution reduces attribution ambiguity and allows direct association with flight trajectory data (e.g., ADS-B) [14]. While their coverage is geographically limited, they provide the most practical basis for building human-labelled contrail-to-flight datasets, such as the GVCCS [8], [9].

A range of algorithmic approaches have been explored for contrail-to-flight attribution. Geometric methods project ADS-B trajectories into image space and match them to detected contrails. [6] combined Mask R-CNN detections

with graph-based assignment strategies, incorporating wind-corrected (i.e., advected) trajectories. [12] employed scalable nearest-neighbour searches with penalties for temporal and spatial mismatch. These methods are interpretable but sensitive to trajectory uncertainties and contrail deformation.

Probabilistic approaches extend geometric matching by explicitly modelling uncertainty. For instance, [15] proposed Bayesian frameworks that account for uncertainties in both contrail geometry and flight tracks, improving robustness. Regional case studies such as [16] attributed MODIS-observed contrails to individual flights, quantifying cirrus coverage and radiative forcing, but relied on indirect physical consistency rather than labelled ground truth.

[13] introduce CoAtSaC, a multi-frame, error-aware approach that attributes each detected (linearized) contrail to a flight by comparing the flight’s advected track to the contrail across multiple satellite frames.

III. DATA


Our contrail-to-flight attribution framework relies on three key types of data: (1) contrails observed by a ground-based camera, (2) flight trajectory data for aircraft passing over the camera, and (3) meteorological measurements at cruising altitudes within the camera’s field of view, used to simulate forecasted contrails. In this section, we describe each of these essential data modalities and explain how they are implemented in our study.

A. Observed Contrails

The primary element of our contrail-to-flight attribution framework is the set of visible contrail instances captured in ground-based camera imagery. To develop and validate our attribution algorithm, we leverage a labelled benchmark that provides human-verified contrail-to-flight pairings. Specifically, we use the GVCCS dataset [8], [9], a curated collection of approximately 24K annotated frames from 122 videos.

Each frame includes multi-polygons delineating visible contrails in pixel coordinates. These annotations capture both newly formed contrails, appearing shortly after formation directly above the camera, and older contrails that originated outside the field of view but drifted into it due to wind advection. Each contrail is manually labelled with a unique identifier and a status tag (“new” or “old”), ensuring temporal consistency across frames and enabling tracking of individual contrails as they evolve. A “new” contrail refers to a contrail that is first observed forming within the camera’s field of view and that can be reliably associated with the flight that generated it. In contrast, an “old” contrail corresponds either to a contrail advected into the camera’s field of view from outside, or to a potentially new contrail that could not be linked to a specific flight (e.g. due to cloud occlusion or other visual limitations). This distinction also enables quantitative evaluation of contrail-to-flight attribution algorithms.

Figure 1 presents a sample image sequence from GVCCS, recorded on 26 April 2024. The GVCCS operates at a temporal resolution of 30 s. For improved clarity and to better illustrate



the contrail evolution, however, Figure 1 displays one frame every 3 minutes. Figure 2 illustrates the human-generated annotations for the same sequence. New contrails are colour-coded and retain consistent identifiers across frames, while old contrails are shown in grey.

B. Flight Trajectory Data

The GVCCS dataset also includes the flight trajectories (ADS-B observations) for all aircraft passing above the camera during each observation sequence. Crucially, these trajectories are linked to the contrail annotations through consistent flight identifiers, enabling direct and rigorous comparison between observed contrails and their originating flights.

For each frame, candidate flight positions are projected into the pixel coordinate system, allowing geometric comparisons with observed contrails. The dataset also includes metadata such as callsigns and aircraft types, which can support filtering (e.g., excluding propeller aircraft) or computing forecasted contrail predictions with physics-based models like CoCiP.

C. Meteorological Data

Reliable contrail-to-flight attribution depends critically on the atmospheric conditions that control a contrail's formation, persistence, and transport. To model this, our method leverages ERA5 reanalysis data from the European Centre for Medium-Range Weather Forecasts (ECMWF). We extract key meteorological fields, including horizontal winds, temperature, pressure, and relative humidity, at isobaric levels between 200 and 300 hPa, encompassing the standard cruising altitudes.

These fields are the essential inputs for simulating contrail advection and evolution. By driving physical models (e.g., CoCiP) or simpler plume or Lagrangian advection models, the ERA5 data allows us to generate a set of plausible contrail trajectories.

For each observed contrail, the framework compares it with simulated contrails generated from candidate flight trajectories in order to identify its most probable source flight. Although we rely on ERA5 for its high spatiotemporal resolution and public availability, the framework is compatible with any gridded meteorological dataset of sufficient quality.

Figure 3 illustrates the forecasted contrails generated by a simple dry advection model for the flights crossing the camera's field of view during the sequence shown in Fig. 1. Dashed lines indicate the actual flight trajectories, while the solid multi-polygons represent the corresponding simulated contrails. Note that when the width of a contrail is very small, the multi-polygons may appear as simple lines.

IV. CONTRAIL-TO-FLIGHT ATTRIBUTION FRAMEWORK

An aircraft emits exhaust continuously along its trajectory. The exhaust released at a given point is transported and deformed by the wind, so an observed contrail at time t generally corresponds to exhaust emitted earlier, at time $t' < t$, and advected to a different location. We formulate contrail-to-flight attribution as the problem of matching observed contrails with theoretical advected contrails generated by flights passing

above the camera. Our framework addresses this task by: (i) restricting comparisons to spatio-temporally plausible pairs, (ii) computing geometric distances, (iii) aggregating distances over time for robustness, (iv) mapping distances to probabilities, and (v) solving the resulting assignment problem.

A. Temporal Filtering

Let c_i denote an observed contrail with formation time t_i , defined as the first time when this contrail identifier was observed. Consider a candidate flight f_j that generates an advected plume (i.e., a forecasted contrail) composed of geometries $\{g_{j1}, \dots, g_{jk}\}$ with corresponding formation times $\{t_{j1}, \dots, t_{jk}\}$. Using CoCiP or a dry advection model¹, for example, each geometry is represented as a simple rectangle, determined by the segment length, width, and orientation.

To retain only the relevant segments of the entire contrail, we introduce two temporal tolerance parameters, Δt_{before} and Δt_{after} , which define the acceptable time window around t_i . A segment g_{jk} is retained if its formation time satisfies:

$$t_i - \Delta t_{\text{before}} \leq t_{jk} \leq t_i + \Delta t_{\text{after}}. \quad (1)$$

If no segments satisfy this condition, the pair (c_i, f_j) is discarded, as the flight is unlikely to have produced the observed contrail. If multiple segments survive, their unary union is computed to form a single composite geometry. This ensures that the full spatial extent of the plume relevant to c_i is considered when computing geometry-based distance metrics, such as overlap or Hausdorff distance, for attribution purposes.


B. Geometric Distance

For each temporally compatible pair (c_i, f_j) , a non-negative geometric distance d_{ij} is computed, where smaller values indicate better agreement between the observed and forecasted contrails. The choice of distance metric depends on the geometry representation of the contrails. For example, if both the observed and forecasted contrails are represented as polygons, the Intersection over Union (IoU) can be used. All distances are computed in the pixel coordinate system. To prevent assigning probability mass to implausible matches, distances exceeding a physically meaningful cut-off are excluded. Specifically, we introduce a threshold τ_d representing the maximum acceptable separation after accounting for advection uncertainty. If $d_{ij} > \tau_d$, we set $d_{ij} = \infty$.

C. Memory-Aware Aggregation

Instantaneous pairwise distances may fluctuate due to intermittent detections or advection mismatches. To stabilize the attribution decisions, a memory mechanism aggregates evidence over time. Specifically, this mechanism maintains a running summary of past distances between each observed and forecasted contrail pair. For example, for each pair (c_i, f_j) , one can compute an exponentially weighted moving average (EWMA) of the instantaneous distance $d_{ij}^{(t)}$ observed at t :

¹<https://py.contrails.org/notebooks/advection.html>






Figure 1. Illustrative sequence of frames on 26 April 2024.




Figure 2. Annotations for the illustrative sequence shown in Fig. 1. The sequence includes two new and one old contrails.




Figure 3. forecasted contrails generated by a simple dry advection model for the flights crossing the camera's field of view during the sequence shown in Fig. 1. Dashed lines indicate the actual flight trajectories, while solid multi-polygons represent the simulated contrails.

$$\tilde{D}_{ij}^{(t)} = \alpha \tilde{D}_{ij}^{(t-1)} + (1 - \alpha) d_{ij}^{(t)}, \quad \alpha \in [0, 1]. \quad (2)$$

Larger values of α emphasize persistence, resulting in stronger smoothing over time, whereas smaller values allow the memory to adapt more quickly to changes in distance.

D. Distances to Probabilities

The next step is to convert geometric distances into probabilities. Given the distance matrix $\tilde{D} \in \mathbb{R}^{m \times n}$ of m contrails and n flights, we first define scores $S_{ij} = -\tilde{D}_{ij}$ for finite entries, assigning a very negative sentinel value to excluded

pairs. Each row is then converted into a probability distribution using a softmax with inverse temperature $\beta > 0$:

$$P_{ij} = \frac{\exp(\beta S_{ij})}{\sum_{k=1}^n \exp(\beta S_{ik})}. \quad (3)$$

To suppress low-confidence associations, a probability floor τ_p is applied: if $P_{ij} < \tau_p$, we set $P_{ij} = 0$ for assignment purposes. The inverse temperature β controls the selectivity of the mapping, with larger values producing more peaked distributions. Note that the probability distribution can be computed either row-wise or over the entire matrix, depending on

the choice of denominator aggregation. This choice primarily depends on the subsequent assignment algorithm.

E. Assignment

Given the probability matrix $P \in [0, 1]^{m \times n}$, the objective is to select the most likely source flight for each observed contrail. This can be done in several ways: independently for each row, or as a one-to-one assignment over the entire matrix using the Hungarian algorithm, for instance.

For each contrail c_i , the framework returns the selected flight identifier (if any), along with the corresponding attribution probability P_{ij} and the aggregated distance $D_{ij}^{(t)}$.

V. PARTICULAR IMPLEMENTATION

In this section we instantiate the general attribution framework introduced above. While the framework is designed to be modular, here we describe the concrete design choices we implemented for evaluation. Each step of the pipeline is specified: how observed and forecasted contrails are represented, how spatial distances are computed and temporally aggregated, and finally how assignments are made.

A. Geometries

The first step is to represent both observed and forecasted contrails in a form suitable for spatial comparison.

1) *Annotations*: Observed contrails are provided as human- or machine-generated annotations in the form of polygon vertices. In our implementation, each polygon is reduced to its central path by applying a morphological thinning operation, which produces a one-pixel-wide structure along the centerline of the shape. The resulting skeleton is represented as a graph of connected pixels, from which we extract the longest path to serve as the contrail geometry.

Figure 4 illustrates the process: (1) the skeleton of the contrail mask is converted into a graph, and (2) the longest path is extracted to represent the contrail geometry. Each polygon is processed independently in the case of multi-polygon annotations.

2) *forecasted contrails*: forecasted contrails are computed using dry advection with an initial width of 100 m. Every 10 seconds, we assume the aircraft emits new exhaust. Each exhaust parcel is then propagated forward using the `pycontrails` dry advection model with a 30-second integration step. This model simulates how the plume moves with the wind (horizontally and vertically) and broadens with time.

At each integration step and for each waypoint, the model outputs a segment characterized by its position, length, orientation, and width, which together define a polygon. Each polygon is also timestamped by its formation time. For temporal filtering, observed contrails are therefore compared only to the union of polygons with compatible formation times.

B. Distance Metric

To compare an observed contrail with a forecasted contrail, we use a directed (asymmetric) variant of the Hausdorff distance. Let A denote the observed contrail geometry (i.e., its




Figure 4. Conversion of an annotated contrail into a geometry.

centerline) and B the union of polygons from the forecasted contrail. The distance from A to B is defined as:

$$d(A \rightarrow B) = \max_{a \in A} \min_{b \in B} \|a - b\|. \quad (4)$$

That is, for each point $a \in A$ we measure the Euclidean distance to the nearest point in B , and we then take the maximum of these nearest distances.

Intuitively, this measures how well the observed contrail is covered by the predicted contrail: if every point of A lies close to some point of B , the distance is small; if even one part of the observation is far from the plume, the distance is large.

Distances exceeding $\tau_d = 30$ pixels are discarded as implausible and replaced with ∞ , preventing spurious matches.

To stabilize the distance signal and improve assignment robustness, we employ an EWMA with $\alpha = 0.7$, balancing persistence with responsiveness to evolving contrail shapes.

C. Assignment Strategy

Finally, we assign each observed contrail to the most likely theoretical source. In this implementation, we use a simple greedy strategy: for each observed contrail, we select the candidate flight with the highest non-zero probability using $\beta = 1$. Each row is processed independently, without enforcing one-to-one constraints. Although this does not guarantee global optimality, it is computationally efficient and works well in practice when contrails have clearly dominant matches.

Because assignment is done per row, we normalize probabilities row-wise, so the sum across candidates for each observed contrail equals one. Furthermore, to increase robustness, we apply a threshold of $\tau_p = 0.5$: only assignments with probability above 50% are accepted. Observed contrails with no confident match are left unassigned to avoid false attributions.

VI. RESULTS

For our experiments, we used the full GVCCS dataset. To isolate the performance of the attribution algorithm from the quality of contrail detection, we treated the provided polygon annotations as perfect detections. This allowed us to focus exclusively on evaluating contrail-to-flight attribution performance. Each newly forming contrail in the dataset is annotated with the flight identifier of the aircraft that generated it. We used this information as ground truth for attribution.

Annotations are categorized as either *new* or *old* contrails. After processing with the contrail-to-flight attribution algorithm, each contrail falls into one of two main categories:

- **Attributed contrail:** the algorithm assigned a flight to the contrail.
- **Unattributed contrail:** no flight was assigned.

For attributed contrails, the outcomes are classified as:

- **Correct Attribution:** the contrail was attributed to the correct flight.
- **Wrong Attribution:** the contrail was attributed to an incorrect flight.
- **False Attribution:** the contrail was actually an old contrail.

For unattributed contrails, the possible outcomes are:

- **Correct Omission:** the contrail was an old contrail.
- **Missed Attribution:** the contrail was new but was not attributed.

This section presents the outcomes of our experimental evaluation. We first provide individual examples to illustrate the possible results of the attribution algorithm, followed by an analysis of its aggregated performance on the full dataset.

A. Illustrative Examples

In the figures below, following the conventions established in previous figures, old contrails are shown in grey. The actual flight paths are indicated with dashed lines, while the forecasted contrails are represented as solid polygons, which may appear as lines when the advected contrail is thin. Annotations are displayed as polygons, even if internally represented as lines by the algorithm, and are coloured according to the attributed flight, if any. These examples highlight the different ways the algorithm can attribute or omit contrails, providing visual insight into its decision-making process.

Figure 5a illustrates a textbook case of correct attribution. Each observed contrail is accurately matched to its generating flight: contrail b9936bed to flight 3453a381, contrail 3452e65e to flight b114466, and contrail 34543e56 to flight 0354c9b5. This example vividly demonstrates the importance of modeling wind advection: strong northwesterly winds shift contrails toward the bottom-right of the frame almost immediately after formation. Without accounting for this drift, aligning contrails to aircraft paths would be challenging.

Figure 5b shows a wrong attribution, highlighting the limitations of the algorithm under ambiguous conditions. Contrail 81a22736 was incorrectly assigned to flight 93aa1dd1, because an advected contrail from a flight that passed through




Figure 5. Illustrative examples of attributed contrails.

the same location seconds earlier coincided closely with the observation. This issue could be mitigated by tightening the temporal filtering on contrail formation, for instance.

Figure 5c shows a false attribution. In this example, an older contrail (941f0955) is incorrectly assigned to flight 336aa26c because the advected path closely aligns with the observed contrail. Such cases illustrate the difficulty of distinguishing persistent contrails from newly formed ones and may also indicate occasional errors in the ground-truth labels.

Figure 6a demonstrates a correct omission. Contrail 6d209a0f drifts into the camera's field of view from outside and does not correspond to any advected forecasted contrails in the frame. The algorithm appropriately leaves it unattributed, avoiding false associations.

Finally, Figure 6b shows a missed attribution under strong easterly winds. A newly formed contrail (0646906f) was not matched to any flight because the temporal and spatial

thresholds prevented a valid candidate from being considered. This example underscores the delicate balance in threshold selection.




Figure 6. Illustrative examples of unattributed contrails.

B. Aggregated Performance

The Sankey diagrams in Figs. 7a and 7b illustrate the performance of the contrail-to-flight attribution algorithm at two distinct points in time: the first frame in which a contrail becomes visible in the camera imagery, and the final frame in which that contrail remains visible. These diagrams provide an intuitive visualization of how contrails are distributed across attribution outcomes, with the width of each flow proportional to the number of contrails that follow that outcome.

At the moment of first appearance (Fig. 7a), the dataset comprises 71.9% new contrails and 28.1% old contrails. The attribution algorithm assigned flights to 73.3% of all contrails, leaving 26.7% unattributed.

Focusing on new contrails, where attribution is crucial, the algorithm correctly identified the generating flight for 67.7% of the dataset. This translates to a striking 94.2% correct attribution rate relative to the new-contrail group. Missed attributions, where new contrails remained unattributed, accounted for only 3.0% of the dataset (4.2% of new contrails), while wrong attributions (i.e., new contrails assigned to the wrong flight) represented just 1.3% of the dataset (1.8% of new contrails). In short, nearly all new contrails were correctly attributed, with only small fractions falling into missed or wrong attribution categories.

The story is different for old contrails. Most old contrails (23.8% of the dataset, or 84.7% of the old-contrail group)




Figure 7. Sankey diagram illustrating the possible algorithm outcomes.

were correctly left unattributed. However, 4.3% of the dataset (15.3% of old contrails) were falsely assigned to a flight. These false attributions may reflect the algorithm's occasional difficulty in distinguishing old contrails from newly forming ones. Although this error affects a relatively small portion of the overall dataset, it is proportionally significant within the old-contrail group and constitutes the main challenge for attribution. Some false attributions may also arise from labelling errors, where contrails labelled as old are in fact new but were not linked to their generating flights by human annotators. Given the inherent complexity of this task, no dataset is perfectly labelled. Carefully reviewing these cases and potentially leveraging the algorithm to identify and correct mislabels could be highly valuable. As a result, the algorithm's true performance for old contrails may be higher than reported.

When performance is examined at the last frame of each contrail's visibility (Fig. 7b), the overall distribution remains largely similar, though some refinements can be observed. Correct attributions rise slightly to 68.0% of the dataset, corresponding to 94.6% of new contrails. This increase is modest in absolute terms but demonstrates that the algorithm becomes marginally more reliable when it has access to the full temporal extent of a contrail. Missed attributions remain constant at 3.0% of the dataset, or 4.2% of new contrails, indicating that additional temporal evidence does not rescue contrails that were initially left unattributed. Wrong attributions, however, decrease from 1.3% to 0.9% of the dataset,

corresponding to a reduction from 1.8% to 1.2% of new contrails. This decline suggests that temporal accumulation helps disambiguate flight candidates for contrails that might have been incorrectly attributed at the moment of first appearance.

The outcomes for old contrails remain unchanged between first and last attribution. This stability indicates that the algorithm's errors with old contrails are not mitigated by longer observation windows and temporal evidence. Once an old contrail is mistakenly treated as new and attributed to a flight, additional frames do not alter that decision. Conversely, if an old contrail is correctly omitted at first appearance, it tends to remain so. These results imply that the challenge of distinguishing between old and new contrails is not one of insufficient temporal context but rather of feature design or classification strategy.

Taken together, the two diagrams provide a coherent overview of the algorithm's performance throughout the contrail lifecycle. In both cases, the algorithm achieves very high attribution accuracy for new contrails. The modest increase in correct attributions and reduction in wrong attributions at the last observation indicate that temporal integration helps disambiguate challenging cases, albeit only slightly. The main limitation of the system is its tendency to falsely attribute a non-negligible fraction of old contrails. Part of this may be due to the inherent complexity of the labelling task: some contrails labelled as old may in fact be new but were not linked to a flight due to occlusion or other visual limitations. The error rate remains roughly constant over time, suggesting that this is a systematic limitation rather than one that can be alleviated by additional temporal evidence.

VII. CONCLUSIONS

This work introduced a modular framework for attributing observed contrails in ground-based imagery to their source flight, based on geometric and temporal reasoning.

From a dataset-level perspective, correct attributions account for roughly two-thirds of all contrails, while correct omissions make up about a quarter. The remaining errors, comprising missed attributions, wrong attributions, and false attributions, collectively represent approximately 9% of the dataset. However, examining performance within the relevant groups highlights the algorithm's strengths. For new contrails, attribution precision reaches around 92–93%, with a recall of approximately 94–95%. For old contrails, correct omissions are frequent but not perfect: about one in six old contrails is falsely attributed to a flight. Some of these errors may reflect genuine algorithm limitations, but a fraction may also stem from labelling errors, where contrails labelled as old are in fact new but were not linked to a generating flight due to occlusion, for instance.

It should be noted that the current framework is limited to clear-sky conditions, as cloud cover or other visual obstructions prevent reliable contrail detection and attribution. Addressing this limitation in future work could broaden the applicability of the system to more diverse atmospheric conditions.

The framework we propose lays the groundwork for several promising extensions. A natural next step is the incorporation of contrail altitude estimation, either directly from imagery or by combining visual data with meteorological or ADS-B-derived information. Altitude cues would help disambiguate overlapping flight paths and improve spatial matching, especially in densely trafficked airspace.

A particularly compelling extension lies in combining ground-based and satellite imagery to track contrails beyond the camera's field of view. Ground cameras offer high-resolution detections and accurate attribution at the early stages of contrail formation. By detecting the same contrail in both ground and satellite imagery, and using the ground-based attribution as a reference, one can establish contrail-to-contrail correspondences across modalities. This enables continued tracking of the contrail's evolution even after it exits the ground camera's view, capturing the full spatial and temporal extent of the contrail. Such an approach could unlock powerful new capabilities in monitoring persistent contrails and assessing their climate impact over time, particularly as they spread into cirrus-like formations far from their origin.

Beyond these immediate improvements, the GVCCS framework opens the door to a range of further research directions. For instance, it could facilitate studies on contrail lifetimes, growth dynamics, and optical properties under different atmospheric conditions.

REFERENCES

- [1] D. S. Lee, D. W. Fahey, A. Skowron, M. R. Allen, U. Burkhardt, Q. Chen, S. J. Doherty, S. Freeman, P. M. Forster, J. Fuglestad et al., "The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018," *Atmospheric Environment*, vol. 244, p. 117834, 2021.
- [2] R. Teoh, Z. Engberg, U. Schumann, C. Voigt, M. Shapiro, S. Rohs, and M. Stettler, "Global aviation contrail climate effects from 2019 to 2021," *EGU sphere*, vol. 2023, pp. 1–32, 2023.
- [3] A. Borella, O. Boucher, K. P. Shine, M. Stettler, K. Tanaka, R. Teoh, and N. Bellouin, "The importance of an informed choice of co 2-equivalence metrics for contrail avoidance," *Atmospheric Chemistry and Physics*, vol. 24, no. 16, pp. 9401–9417, 2024.
- [4] U. Schumann, "A contrail cirrus prediction model," *Geoscientific Model Development*, vol. 5, no. 3, pp. 543–580, 2012. [Online]. Available: <https://gmd.copernicus.org/articles/5/543/2012/>
- [5] K. Gierens, S. Matthes, and S. Rohs, "How well can persistent contrails be predicted?" *Aerospace*, vol. 7, no. 12, 2020.
- [6] R. Chevallier, M. Shapiro, Z. Engberg, M. Soler, and D. Delahaye, "Linear contrails detection, tracking and matching with aircraft using geostationary satellite and air traffic data," *Aerospace*, vol. 10, no. 7, p. 578, 2023.
- [7] J. Van Huffel, R. Ehrmanntraut, and D. Croes, "Contrail detection and classification using computer vision with ground-based cameras," in *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE, 2025, pp. 1–6.
- [8] G. Jarry, P. Very, F. Ballerini, and R. Dalmau, "GVCCS: Ground visible camera contrail sequences," Jul. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.15743988>
- [9] G. Jarry, R. Dalmau, P. Very, F. Ballerini, and S.-D. Bocu, "GVCCS: A dataset for contrail identification and tracking on visible whole sky camera sequences," *Earth System Science Data Discussions*, vol. 2025, pp. 1–30, 2025.
- [10] T. M. Fritz, S. D. Eastham, R. L. Speth, and S. R. Barrett, "The role of plume-scale processes in long-term impacts of aircraft emissions," *Atmospheric Chemistry and Physics*, vol. 20, no. 9, pp. 5697–5727, 2020.



- [11] J. Itcovitz, A. Sarna, M. Stettler, and V. Meijer, "Flight-matched contrail observations from GOES - insights from aircraft-engine analysis," in *AGU Fall Meeting Abstracts*, vol. 2024, Dec. 2024, pp. A51M–1867.
- [12] S. Geraedts, E. Brand, T. Dean, S. Eastham, C. Elkin, Z. Engberg, U. Hager, I. Langmore, K. McCloskey, J. Ng, J. Platt, T. Sankar, A. Sarna, M. Shapiro, and N. Goyal, "A scalable system to measure contrail formation on a per-flight basis," *Environmental Research Communications*, vol. 6, 01 2024.
- [13] A. Sarna, V. Meijer, R. Chevallier, A. Duncan, K. McConnaughay, S. Geraedts, and K. McCloskey, "Benchmarking and improving algorithms for attributing satellite-observed contrails to flights," *EGUsphere*, vol. 2025, pp. 1–58, 2025.
- [14] J. Low, R. Teoh, J. Ponsonby, E. Gryspeerdt, M. Shapiro, and M. E. Stettler, "Ground-based contrail observations: comparisons with reanalysis weather data and contrail model simulations," *Atmospheric Measurement Techniques*, vol. 18, no. 1, pp. 37–56, 2025.
- [15] E. Riggi-Carolo, T. Dubot, C. Sarrazat, and J. Bedouet, "AI-driven identification of contrail sources: Integrating satellite observation and air traffic data," *Journal of Open Aviation Science*, vol. 1, no. 2, 2023.
- [16] D. P. Duda, P. Minnis, L. Nguyen, and R. Palikonda, "A case study of the development of contrail clusters over the great lakes," *Journal of the Atmospheric Sciences*, vol. 61, no. 10, pp. 1132 – 1146, 2004.