



AI- AND DATA-DRIVEN IDENTIFICATION OF CONTRAIL SOURCES

Simon Kröger, Sebastian Lück , Jan Goeing 

Technische Universität Braunschweig
Institut für Flugantriebe und Strömungsmaschinen,
Hermann-Blenk-Str. 37, 38108 Braunschweig, Germany

Abstract

Contrails significantly contribute to aviation-induced global warming. Accurate prediction and simulation models of contrails are essential both to better estimate their climate impact, which remains subject to significant uncertainties, and to predict when and where flights will produce contrails, with the long-term goal of rerouting them to reduce their impact. This study investigates whether automated detection of contrails in geostationary satellite imagery, combined with the attribution to the flights that produced them, can create a database of real-world observations capable of validating contrail prediction models. We used an open-source neural network to detect contrails in geostationary satellite images. To match these contrails with flights, we developed a two-stage matching algorithm, building on and extending previous work in this field. The first stage applies geometric filters to narrow down the number of candidate flights per contrail, while the second stage uses wind data to calculate the predicted contrail track for each previously matched flight trajectory and compares it with the observed contrail locations. Additionally, a scoring method was introduced to quantify match quality, distinguishing well-aligned from poorly aligned associations. The system also enables the retrieval of environmental conditions present at the time of contrail formation for each identified match. Compared to an existing matching algorithm, our approach produced one-third fewer matches per contrail on the same dataset, and the individual results appear more plausible. A statistical evaluation shows that while a single contrail cannot always be attributed to one specific flight, the method consistently narrows the results to a small set of candidate flights per contrail. A case study shows that manually reviewing the results, particularly by combining multiple observations of the same contrail over time, helps to further narrow down the list of potential source flights. This demonstrates that although the current version cannot yet automatically identify a contrail's source flight, further improvements could make this possible. In particular, automatic contrail tracking combined with dedicated filtering logic and refined contrail detection may enable the creation of a reliable large dataset of contrail-flight associations usable for model validation.

Keywords

contrails; matching contrails with aircraft; contrail detection; aviation climate impact; satellite imagery

1. INTRODUCTION

Contrails (condensation trails) are ice clouds that form when hot, humid exhaust gas from aircraft engines mixes with cold ambient air, causing water vapor to condense and freeze into ice crystals. If these ice clouds form in ice-supersaturated regions (ISSRs), they do not dissolve within minutes but can persist for hours and grow into extensive cirrus clouds [1–3]. These clouds reflect thermal radiation, thereby influencing the global energy balance, with a net warming effect typically expressed as an increase in radiative forcing [4, 5].

Despite uncertainties, contrails are considered one of the most significant climate impacts of aviation, alongside nitrogen oxides (NO_x) and carbon dioxide (CO_2), each contributing roughly one third of aviation's total climate effect [4–6]. As contrails only form under specific atmospheric conditions, a relatively small

number of flights are responsible for a large portion of this effect. Identifying such flights in advance and potentially rerouting them could therefore significantly reduce the contrail induced climate impact [7, 8]. Beyond their role in operational mitigation, contrail simulation models are also essential for assessing the climate impact of contrails, which remains subject to large uncertainties [4, 5]. Several models have been developed to simulate and/or predict contrail formation, but the lack of a comprehensive validation dataset means their accuracy is largely unknown, limiting confidence in their predictions [5, 9]. This study investigates whether automated detection of contrails in geostationary satellite imagery, combined with the attribution to the flights that produced them, can create a database of real-world observations capable of validating contrail prediction/simulation models.

The task of identifying and linking contrails to flight movements can be structured into four main steps:

- 1) **Data acquisition:** Collecting satellite imagery, air traffic data, and environmental data such as wind, temperature, and humidity.
- 2) **Contrail Detection:** Identifying contrails in the imagery using an AI-based detection model.
- 3) **Matching:** Associating flight trajectories with detected contrails using a filtering logic.
- 4) **Evaluation:** Presenting results to support quality assessment.

The first two steps of this workflow have become significantly easier in recent years due to improved data availability and advances in AI-based image recognition. Most notably, Ng *et al.* recently released OpenContrail, a large labeled dataset for AI-based contrail detection [10], which has been a key enabler, allowing researchers to identify contrails more reliably in satellite imagery. Building on these advances in detection, this study focuses primarily on the matching task: formatting and structuring the data and model outputs, developing a logic to associate them and presenting the results. Two studies are particularly noteworthy in this context, as they have already explored this problem: Riggi *et al.* [11] and Chevallier *et al.* [12].

Riggi *et al.* [11] simplified contrails and flights as lines defined by two parameters, ρ and θ (similar to the slope and y-intercept of a linear function). By applying tolerances to these parameters, they created a rectangle in ρ - θ space for each contrail. A flight is considered a match if its parameters fall within this rectangle. In addition, the method checks whether the flight is moving toward or away from the contrail and whether the environmental conditions along the trajectory satisfy the thermodynamic requirements for contrail formation to further narrow down the number of flights associated with a contrail. The general approach of simplifying contrails and flight trajectories and filtering based on their properties has been adopted in this study, although in a different form (see Section 2.4), as has the approach of filtering by relative flight direction (see Section 2.3.3).

Chevallier *et al.* [12] first filtered flights based on their proximity and orientation relative to the contrail. In a second step, they calculated predicted contrail tracks (PCTs) by determining where a contrail caused by a specific aircraft would be located at the time of observation based on wind data. These PCTs were then compared with the actual contrail observations from satellite imagery, the observed contrails track (OCT). They also implemented a tracking algorithm to automatically follow contrails across multiple satellite images over time, enabling comparison of the motion of PCTs with that of the OCTs. We adopt the general concept of comparing PCTs with OCTs, although the filtering logic in this study differs. The contrail tracking approach has not

been implemented but represents a logical extension of our method for future work (see Section 3.3).

The method developed in this study builds on both approaches while introducing new concepts to improve matching accuracy. Most notably, the filtering logic was redesigned to include a new form of OCT-PCT comparison and a novel proximity-based filtering method. In addition, a scoring metric was implemented to quantify match quality. Together with a backtracking procedure that identifies the section of a flight trajectory responsible for a contrail formation, which enables the investigation of environmental conditions at the time of formation.

2. METHOD

2.1. Contrail Matching Workflow

Figure 1 illustrates the overall workflow for matching contrails with flights, organized into the four steps outlined in Section 1. First, the required data (traffic data, satellite imagery, and meteorological data) is collected (1). The satellite image is then preprocessed to meet the model's input requirements (2a), the model is applied (2b), and the resulting output is used to extract individual contrail lines for matching purposes (2c). Next, the traffic data is preprocessed by removing invalid entries, filtering out low-altitude data points, and segmenting flight trajectories into linear elements (3a). In the first filtering stage, geometric filters are applied based on relative flight direction (direction filter), geographic proximity (location filter), and orientation (orientation filter), reducing the number of candidate flights and establishing an initial match between contrails and flights (3b). In the second filtering stage, for each flight retained from the first stage, a predicted contrail track (PCT) is calculated and compared with the observed contrail track (OCT) to further refine the matches (3c). In the final step, a match quality score is calculated for each match to distinguish well-aligned from poorly aligned PCT-OCT pairs (4a). Additionally, the time and location of contrail formation are traced back to determine the environmental conditions at the point of formation (4b).

2.1.1. Satellite Images (1a)

Satellite imagery was sourced from the geostationary GOES-16 satellite, positioned at 75.2°W, with a temporal resolution of 10 minutes and a spatial resolution of 2x2 km per pixel in the infrared spectrum. The main advantages of these images are their high temporal resolution and large coverage area. The primary disadvantage is the relatively coarse spatial resolution due to the satellite's distance from Earth, which prevents direct observation of aircraft and newly formed contrails. Consequently, there is always a gap of anywhere from a few minutes to half an hour [12] between the time a contrail is emitted and when it is first ob-

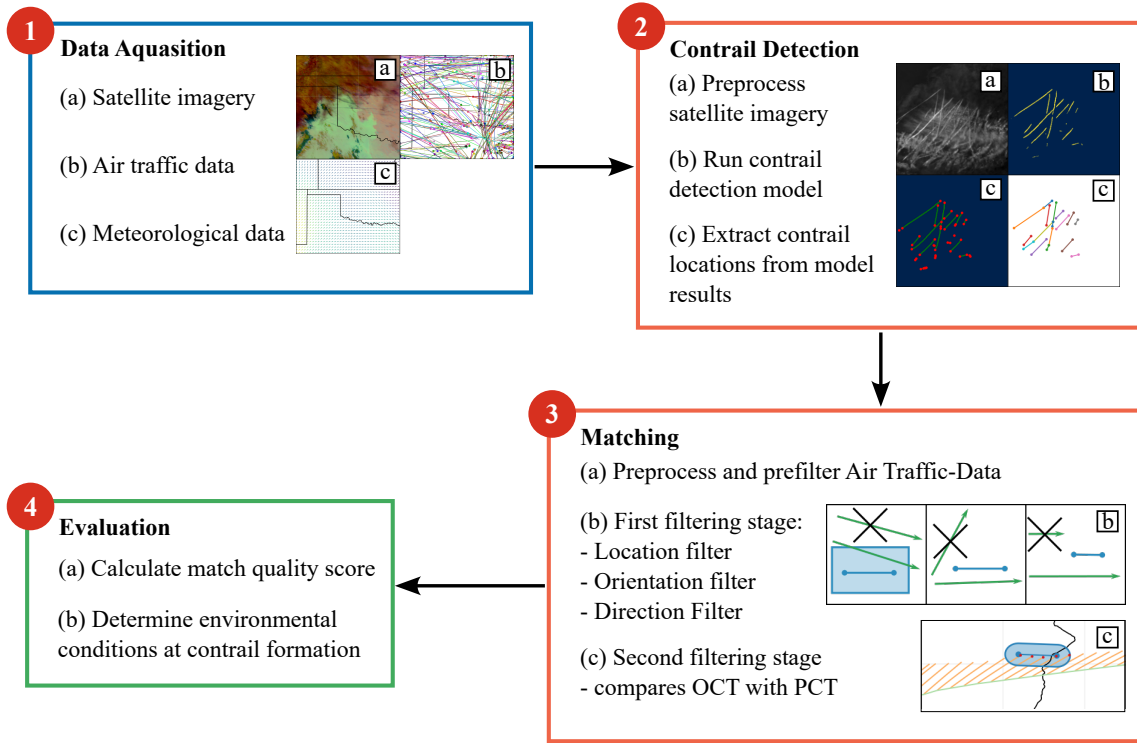


FIG 1. Workflow of contrail detection and matching used in this study

served. This gap represents the core challenge of this study, as the problem would become trivial if the emitting flight could be directly observed.

2.1.2. Air Traffic Data (1b)

Air traffic data is obtained from the OpenSky Network [13]. The data is accessed and processed via the *traffic* Python library [14], resampled to 10-second intervals, and processed using the built-in `clean_invalid()` function and a consistency filter. Points along the flight trajectories below 8 km were excluded, as contrail formation is unlikely at these altitudes. According to the Schmidt–Appleman criterion (SAC) and the NASA Standard Atmosphere, contrails cannot form below about 8.4 km [3].

2.1.3. Meteorological Data (1c)

Meteorological data used to calculate wind displacement and retrieve environmental conditions at contrail formation are obtained from the ERA5 reanalysis dataset provided by ECMWF [15]. The data are published on a regular latitude–longitude grid with a cell size of $0.25^\circ \times 0.25^\circ$ and 37 pressure levels from 1 hPa to 1000 hPa, with a temporal resolution of one hour. For the evaluation of relative humidity with respect to ice (RHi), known deviations were corrected using the method by Teoh et al. [16, 17]. Data access and processing were carried out using the Python library *pycontrails* [18].

2.2. Contrail Detection (2)

A publicly available contrail detection model developed by Junzi Sun [19] was used in this study. The model is based on a ResUNet architecture and processes 320×320 pixel single-channel greyscale image. The model was not retrained, and the original weights provided by the author were used. After testing, the *own-focal-1000epoch* weights were selected because they provided the most conservative performance and minimized false positives. Missing a contrail occasionally was considered less problematic than detecting one where none existed. The model outputs pixel-wise probabilities for contrail presence, and a 60% probability threshold was applied to classify pixels as part of a contrail.

Before running the model, the satellite images were preprocessed to meet its input requirements (2a). First, the images were projected onto a map using a Plate-Carrée projection, to allow the extraction of specific regions by latitude and longitude. A region of $5^\circ \times 5^\circ$ was typically selected for this study. The images were then converted into Brightness Temperature Difference (BTD) grayscale images and scaled to 320×320 pixels (as shown in Figure 2a). If the extracted region was not rectangular, black padding was added to achieve the required shape. The BTD is defined as the normalized temperature difference between GOES-16 wavelength channel 13 ($10.3 \mu\text{m}$) and channel 15 ($12.3 \mu\text{m}$). The resulting image is then fed into the contrail detection model, which returns a binary 2D array (as shown in Figure 2b) indicating for each pixel whether a contrail is present (2b).

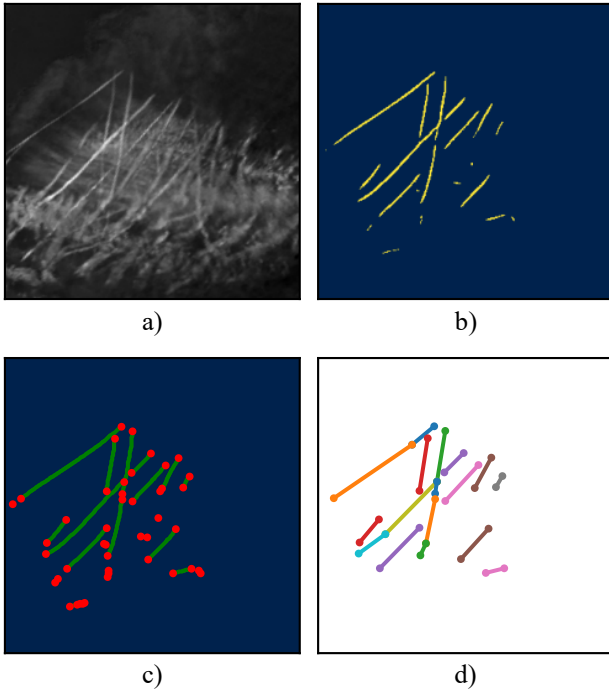


FIG 2. Contrail detection workflow: (a) preprocessing of satellite images into Brightness Temperature Difference (BTD) format, (b) binary output of the AI-based model, (c) detected structures by the skeleton-network algorithm, and (d) extracted contrail line segments.

The filtering algorithm used in this study is based on comparing linear segments. Therefore, the output of the contrail detection model must be converted into straight lines with geographical coordinates (2c). The procedure is illustrated in Figure 2 c and d. First, a so called skeletonization algorithm from the `scikit-image` library is applied, producing a one-pixel-wide centerline of the structures detected by the contrail detection model (see Figure 2c, green lines). Next, a skeleton-network algorithm [20] is used to identify endpoints and intersection points (see Figure 2 c, red points) within the network of lines. This information is then used to define individual lines, which are simplified using a Douglas-Peucker algorithm. Only segments exceeding a minimum length threshold of 10 pixels (corresponding to about 17–18 km in the $5^\circ \times 5^\circ$ areas used in this study) are retained as contrails (see Figure 2d). Finally, these lines are transformed from pixel-coordinates to geographical-coordinates and used in the following filtering algorithms.

2.2.1. Flight Trajectory Segmentation (3a)

For the filtering method used here, straight segments of the flight trajectories are required. Therefore, after the prefiltering and preprocessing described in Section 2.1.2, the trajectories are simplified using the Douglas-Peucker algorithm (as shown in Figure 3). This reduces the number of points by retaining only

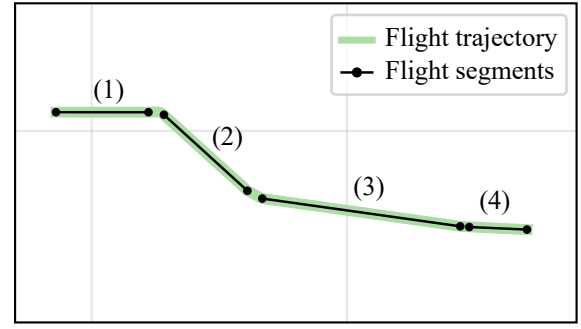


FIG 3. Example of the flight trajectory segmentation.

the start and end points of straight sections. The maximum permitted distance between the simplified and original line is set to 100 m, which provides a good compromise between minimizing the number of segments and maintaining sufficient accuracy. Because this simplification preserves more points in curved sections than in straight sections, only segments exceeding a minimum length of 10 km are retained to reduce the number of segments. The algorithm was implemented using the `simplify()` function from the `traffic` library [14].

2.3. First Filtering Stage (3b)

The goal of the this first filtering stage is twofold: first, to establish an initial match between contrails and flights, and second, to remove flights that clearly could not have been the source of the contrail. Inspired by the work of Riggi et al. [11], the comparison is performed using lines. Unlike in their approach, however, lines with defined endpoints are used instead of infinite lines in order to improve accuracy. Therefore, as explained in Sections 2.2 and 2.2.1, flights and contrails are first split into line segments before applying the filter. Each segment is then evaluated against three criteria: 1) geographic proximity to the contrail (location filter), 2) alignment with the contrail (orientation filter), and 3) movement relative to the contrail (direction filter) (see Figure 4). If a segment passes all three criteria, the corresponding flight is considered an initial match.

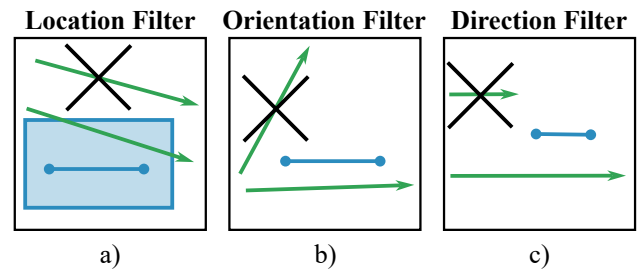


FIG 4. Schematic illustration of the first filtering stage. Flights in green, contrails in blue, and crossed-out flights indicate those discarded.

Because a subsequent filtering step follows, all thresholds in this stage are chosen conservatively to allow slightly more potential matches rather than risk missing one. This method significantly reduces the number of possible flights per contrail, which decreases the load on the following computationally more demanding second filtering stage.

2.3.1. Location Filter

The location filter (see Figure 4 a) evaluates the geographic proximity of a flight trajectory segment by checking whether it intersects a rectangular area surrounding the contrail (see blue rectangle in Fig. 4a and Fig. 5). This area is determined by backtracking the wind displacements of the contrail endpoints to the start time of the considered traffic data. Since the contrail altitude is unknown, this is done for every considered altitude level of the meteorological dataset, resulting in multiple wind displacements for each contrail endpoint (see orange lines in Figure 5). From the resulting displacements, the minimum and maximum latitudes and longitudes values are selected to form a bounding rectangle, which is then expanded by a tolerance to account for uncertainties in the calculations. Figure 5 illustrates this process for a two-hour backtracking period with a tolerance of 15 km. The wind in this example is from the northwest. Wind displacement is calculated according to the method described in Section 2.4, but using a negative time step to go backwards in time. Intersection between a flight trajectory segment and the contrails location filter rectangle is determined using the `intersects()` function from the Python *shapely* library.

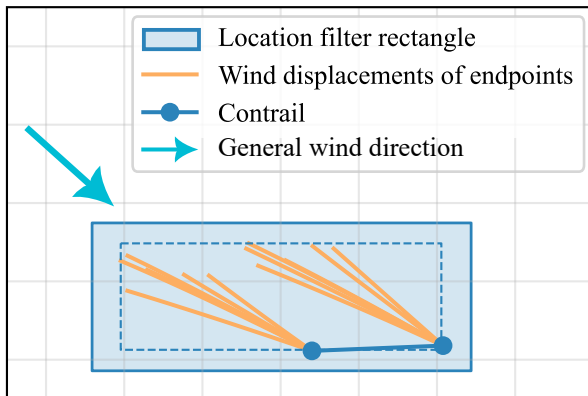


FIG 5. Construction of the location-filter rectangle. A flight segment is retained if it intersects this rectangle (see blue rectangle in Figure 4 a).

2.3.2. Orientation Filter

The orientation filter (Figure 4 b) checks whether a flight trajectory segment is aligned in a way that it could have produced the contrail. Two orientation ranges are defined around the contrail heading to account for both possible flight directions. Segments out-

side these ranges are discarded. A constant tolerance of $\pm 20^\circ$ is applied to compensate for uncertainties in the data, contrail detection, and potential rotation of the contrail due to wind. At typical aircraft cruise speed and assuming a constant wind orthogonal to the flight direction, this tolerance corresponds to a maximum wind speed of approximately 300 km/h, meaning that winds above this threshold could cause true contrail-producing flights to be incorrectly filtered out. Riggi et al. [11] chose a tolerance for the orientation parameter (in their case θ) of $\pm 3^\circ$, and Chevallier et al. [12] used a range of $\pm 7^\circ$. However, considering the implication above and experience from testing, these values can be too low in some cases.

2.3.3. Direction filter

The direction filter checks whether a flight trajectory segment is moving toward or away from the contrail (Figure 4 c). A segment moving toward the contrail cannot have caused it, as the flight would not yet have been in the contrail's location. This method, adapted from Riggi et al. [11], measures the distances between the endpoint of a flight trajectory and the endpoints of the contrail. If these distances are greater than those from a slightly earlier point along the trajectory, the flight is assumed to be moving away from the contrail. If both distances increase, the segment passes the filter. Otherwise it is discarded.

2.4. Second Filtering Stage (3c)

In the second filtering stage, for every flight identified as a potential contrail source in the first filtering stage, the wind displacement of geographical points along the flight trajectory is calculated from the time the aircraft passed each point until the time of the contrail observation to create the Predicted Contrail Track (PCT, see red line in Figure 6). This PCT is then compared to the observed contrail track (OCT), extracted from the satellite image by the contrail detection algorithm (see Section 2.2), to further reduce the number of potential candidate flights per contrail.

The wind displacements are computed using the explicit Euler method. The chosen time step of 60s is well below the instability limit defined by the CFL condition ($CFL \ll 1$) [21] and, based on testing with various time steps, offers a balance between computational efficiency and accuracy. For other spatial resolutions of meteorological data, this value should be re-evaluated, as Euler method stability depends on the relationship between spatial resolution and the maximum distance traveled between time steps. Changes in contrail altitude, such as those caused by downwash, are not considered in the calculation.

The filtering process uses the PCT to evaluate both its spatial and directional agreement with the OCT. Spatial agreement is assessed by checking whether each point of the PCT lies within a 15 km tolerance zone around the observed contrail (purple

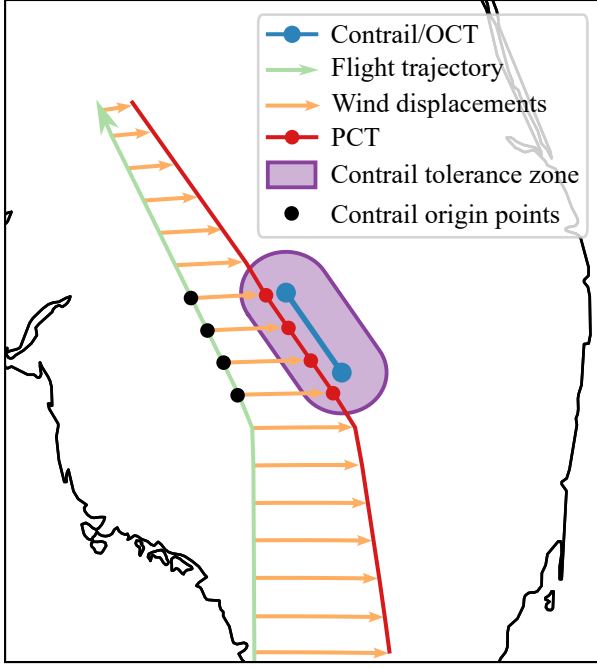


FIG 6. Illustration of the PCT–OCT comparison used in the second filtering stage. PCT points within the contrail tolerance zone are marked in red, and their corresponding origin points, used in Section 2.5 to determine the environmental conditions at formation, are shown in black.

shaded area in Figure 6). If more than two such points (see red points in Figure 6) fall inside the tolerance zone, the contrail is retained. Directional agreement is then evaluated by drawing a line between the first and last PCT points located within the contrail tolerance zone (i.e., the highest and lowest red points in Figure 6). The orientation of this line must be within a $\pm 10^\circ$ tolerance of the OCT heading, similar to the orientation filter described in Section 2.3.2. Flights that meet both criteria are retained, while all others are discarded.

2.5. Determining Environmental Conditions at Contrail Formation (4b)

Based on the calculations performed in the second filtering stage (see Section 2.4 and Figure 6), the time and location of contrail formation can be determined by backtracking the wind displacements (orange vectors in Figure 6) to their origin (black points in Figure 6). This allows the time, location, and altitude of contrail formation to be identified, and the environmental conditions at the time of formation to be derived by interpolating the ERA5 meteorological data.

2.6. Match quality score (3a)

To distinguish well-aligned PCT–OCT pairs from those barely within the tolerances defined by the second filtering stage, a custom score was developed, as no consistent quality metric for this application could

be found in the literature. The total score $s_t \in [0, 1]$ is defined as a combination of three components, each represented by its own subscore: proximity (s_p), orientation (s_o), and length (s_l), and is calculated as

$$(1) \quad s_t = s_l \cdot \left(\frac{s_o + s_p}{2} \right).$$

The proximity score s_p is calculated using Equation 2, where d_{\max} is the distance between the contrail and the edge of the tolerance zone (here 15 km), and d_{mean} is the mean distance between the PCT points within the contrail tolerance zone (red points in Figure 6) and the contrail itself.

$$(2) \quad s_p = 1 - \frac{d_{\text{mean}}}{d_{\max}}, \quad d_{\text{mean}} = \frac{1}{n} \sum_{i=1}^n |d_i|.$$

The orientation score s_o is computed from the angular difference $\Delta\beta$ between the observed contrail heading β_c and the heading of a line through the first and last in-zone points β_f , with a maximum allowed deviation $\beta_{\max} = 10^\circ$:

$$(3) \quad s_o = 1 - \frac{\Delta\beta}{\beta_{\max}},$$

$$(4) \quad \Delta\beta = \min(|\beta_c - \beta_f|, 360^\circ - |\beta_c - \beta_f|).$$

The length score s_l is defined as the ratio of the PCT length within the contrail tolerance zone (l_f , see red points in Fig. 6) to the OCT length (l_c), capped at 1.

$$(5) \quad s_l = \begin{cases} \frac{l_f}{l_c}, & l_f < l_c \\ 1, & l_f \geq l_c. \end{cases}$$

Figure 7 shows three examples. In (a) all subscores are high, resulting in a near-perfect total score s_t . In (b) orientation s_o and proximity s_p alone would give $s_t = 0.39$, but the short matched segment length reduces the total score to 0.23. In (c) orientation and length match well, but poor proximity significantly lowers the total score.

3. RESULTS

The performance and accuracy of the developed method are evaluated using three approaches. First, the results are compared with those of a similar method from the literature by applying both to the same datasets and examining the differences. Second, the method is tested across various times and locations to assess its quantitative performance. Third, in a case study, the algorithm is applied to the same contrail observed at different times to evaluate its consistency and analyze its behavior.

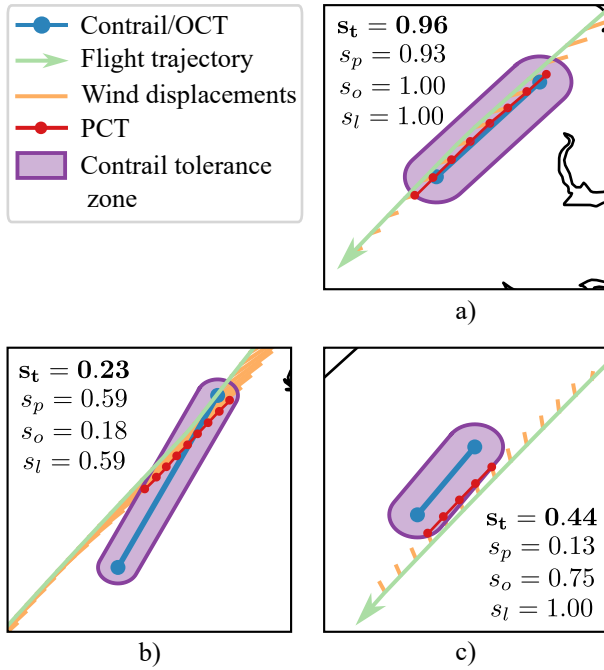


FIG 7. Three examples with their corresponding scores: (a) a near-perfect score, (b) average proximity with a low orientation score, and (c) low proximity with decent orientation.

3.1. Comparison with Riggi et al.

Riggi et al. [11] have made their example dataset and results publicly available in a Git repository. We therefore used it for benchmarking the method presented in this study. The dataset consists of a GOES-16 satellite image from 28 April 2019 at 19:00 UTC, covering the area from 119.65°W to 114.29°W and 30.83°N to 35.37°N (US West Coast), along with corresponding air traffic data from 18:20 UTC to 19:00 UTC for a slightly larger surrounding area. When applied to the same dataset, the average number of matches identified per contrail is about one third of the number found by Riggi et al., even though initially more flights were considered during filtering due to the lower minimum altitude threshold (8 km instead of 10.36 km) and the inclusion of curved flight trajectories.

Figure 8 compares the flights matched by our method with those matched by Riggi et al. for the same contrail. In this case, all flights matched by our method are a subset of those matched by Riggi et al. The prevailing wind direction is from the south-west, which explains why all matched flights are located to the left of the contrail. With our method, only flights that actually passed the contrail were matched, rather than flights that were merely located to the left and similarly oriented. This results from [11] representing flights as infinite lines with two parameters (ρ and θ) instead of finite line segments defined by endpoints used in our method. Furthermore, some flights may indeed have passed the contrail, but the corresponding

data points are not included in the traffic dataset, which covers only a 40-minute timeframe. This indicates that the chosen timeframe may be too short for our method, as it requires a dataset covering the actual contrail formation time to successfully match flights and contrails.




FIG 8. First comparison of the matched flight trajectories between our method and that of Riggi et al.

In Figure 9 another comparison of the matched flights is shown for a different contrail. In this case, the two methods matched completely different flights. Our method did not match flights farther away because the wind displacement calculated by the second filtering stage was insufficient to reach the observed contrail position. Instead, it matched flights closer to the contrail that were discarded by Riggi et al. due to their orientation differing too much from that of the contrail. This shows that the orientation tolerance of $\theta \pm 3^\circ$ might not be sufficient, as the orientation of a contrail can deviate from that of its originating flight more than 3° , particularly under crosswind conditions (see Section 2.3.2).

3.2. Statistical Evaluation

To assess the overall performance of the developed method, a statistical evaluation was conducted using 50 distinct contrail observation situations between January 1 and June 30, 2024. Each situation consists of a $5^\circ \times 5^\circ$ satellite image over a different location in North America, selected at evenly spaced time intervals, with each location manually chosen to contain at least one contrail (see rectangles in Figure 10). The corresponding air traffic data covers an extended area of 200 km in all directions around the satellite image area and includes traffic data up to two hours prior to the satellite image. After filtering flights below 8 km altitude, the dataset contains a total of 36,091 flights, corresponding to an average of 720 flights per situation. In total, 526 contrails were detected (average: 10.5 per situation). As described previously,

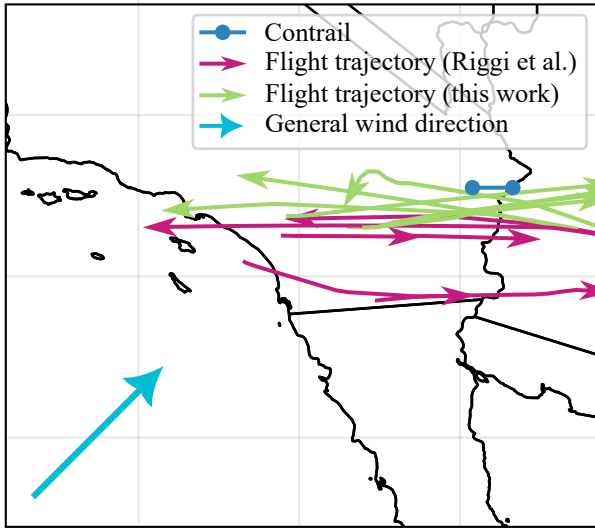


FIG 9. Second comparison of the matched flight trajectories between our method and that of Riggi et al.

a contrail is often segmented into multiple parts and thus counted more than once (see Figure 2.2d).




FIG 10. Geographical areas used in the statistical evaluation.

On average, 2.72 matches were found per contrail, with about 30% of contrails having no matches at all (Figure 11 shows the frequency of different numbers of matches). Possible reasons include false detections by the contrail detection model, a temporal window for the traffic data that does not extend far enough into the past, or inaccuracies in the wind data that, over time, lead to deviations large enough for the PCT to fall outside the contrail tolerance zone. The distribution of the match quality score is roughly balanced, with 44% of matches scoring below 0.5 and 56% above, and about 4% achieving a score greater than 0.9.

To determine whether a significant portion of the flights identified as contrail-causing are actual originators rather than a random selection, the environmental conditions at contrail formation (see Section 2.5) were compared with those of all other flight data points in the statistical analysis. Specifically, we compared a dataset containing every data point from all flights with a subset containing only the data points identified as the source of a contrail (black points in

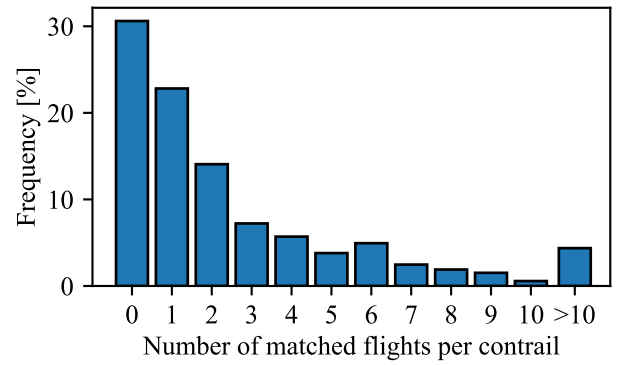


FIG 11. Distribution of the number of matches per contrail in the statistical evaluation.

Figure 6). Furthermore, an additional contrail-causing dataset was created, including only matches with a high score (>0.7), to evaluate whether a high score indeed corresponds to a higher likelihood of being the actual source of a contrail. This comparison is expected to show that the environmental conditions associated with contrail-causing flights more closely align with thermodynamic predictions, most notably satisfying the Schmidt–Appleman criterion and being ice-supersaturated.

When interpreting the results of this analysis, it should be noted that the method may be prone to bias, as satellite images containing many contrails are likely to cover larger areas of contrail-favourable conditions. More contrails lead to more matches, which in turn increases the likelihood of identifying contrail-favourable conditions. To mitigate this, only satellite images containing at least one contrail were chosen. Furthermore, the comparison between the contrail-causing subset and the high-score contrail-causing subset should not be affected by this bias.

TAB 1. Share of data points where the Schmidt–Appleman criterion (SAC) is satisfied and those located in an ice-supersaturated region (ISSR).

	SAC (%)	ISSR (%)
All flight data	69.33	16.24
Subset identified as contrail-causing	75.36	27.08
Subset identified as contrail-causing (score > 0.7)	80.59	26.50

In Table 1, the Schmidt–Appleman criterion (with an engine thermal efficiency = 30%) and ice-supersaturation checks were applied to all three datasets. The table shows the percentages of cases where these conditions were met. As shown, the subset identified as contrail-causing has a higher percentage of cases meeting the Schmidt–Appleman criterion and a higher percentage of ice-supersaturated conditions. The same applies to the high-score contrail-causing

subset, although no difference in ISSR occurrence is observed between the contrail-causing and high-score contrail-causing subsets. To examine this behavior in more detail, the distributions of temperature and humidity values across the different datasets were compared. This was done by defining evenly spaced intervals (10% RH_i for humidity, 3 K for temperature) and plotting the percentage of values within each interval as bars in Figure 12. When comparing RH_i values between the contrail-causing and high-score contrail-causing datasets, the proportion of values above 100% RH_i is very similar, but the proportion of values just below 100% RH_i has significantly increased. There is also no difference in this behavior when using the uncorrected data (see Section 2.1.3). This may be due to uncertainties in relative humidity data, suggesting that a filter based on the SAC and the meteorological dataset used here may not be sufficiently accurate. Chevallier et al. [12] reported similar experiences when implementing a filter based on the SAC.

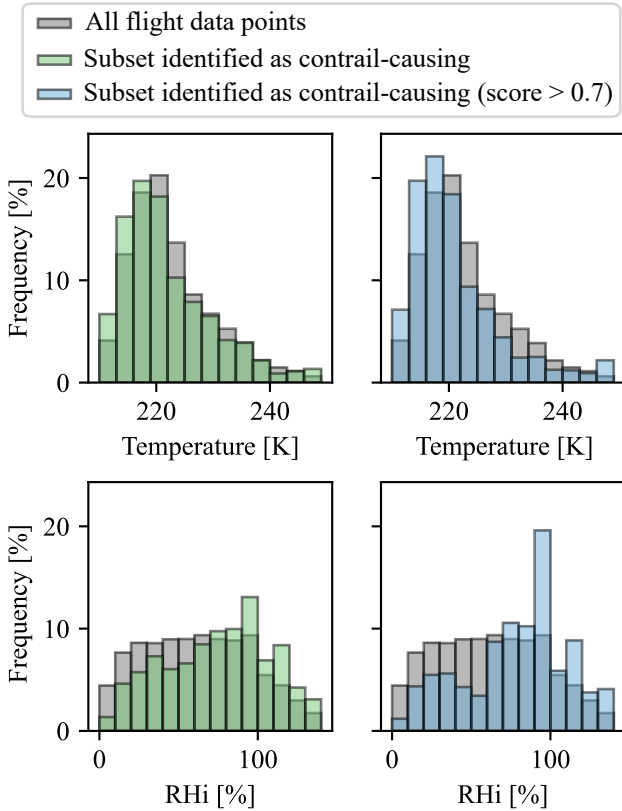


FIG 12. Comparison of the distributions of temperature and relative humidity with respect to ice (RH_i) across the different datasets.

Overall, the temperature and humidity distributions tend toward lower temperatures and higher humidity, and this effect is more pronounced with a higher match quality score, as predicted by the SAC. This indicates that the matched flights are not entirely random but include a significant number of actual contrail originators, and that a higher match quality score corre-

sponds to a higher probability of being the source of the contrail.

3.3. Case Study: Contrail Tracking

The objective of this section is to evaluate the robustness of the developed method by tracking a contrail across multiple satellite images taken at different times and assessing whether it consistently identifies the same flight as its source. Since a contrail can only originate from a single flight, the algorithm should ideally return the same source regardless of the observation time.

The area used in this example, located between 30°N and 35°N and between 97°W and 92°W, was observed on February 10, 2025. Corresponding traffic data were retrieved for a slightly larger bounding box and for the time period from 13:00 UTC to 16:00 UTC. After applying the altitude pre-filter (workflow step 3a), 1,013 flights remained for analysis.

The chosen contrail first becomes visible in the 15:02 UTC satellite image and remains identifiable in subsequent images taken at 10-minute intervals until 15:42 UTC, when it has dissipated to the point that its linear structure is no longer clearly distinguishable. The linking of contrails between consecutive images (i.e., identifying that a contrail in one image is the same as in the next) was performed manually, as was the association of multiple linear contrail segments within a single image to represent one continuous contrail.

In total, 18 different flights were matched with at least one contrail segment at one of the observation times. On average, a single segment at a given time had 5.2 matches. Only three flights matched more than 50% of the segments, and only one flight matched every segment. Many flights with only one to three matches could be ruled out, as their emission time occurred after the first observation of the contrail.

Among the two most likely candidates with the highest geographical scores and the most matched segments, both followed nearly the same route at almost the same time but at different altitudes. One flight flew significantly higher, and thus at colder temperatures, making it the more probable source of the contrail. This flight also had a higher average match quality score and matched every contrail segment rather than all but one. The PCTs and OCTs for this flight are shown in Figure 13, where they generally align well. The larger deviations at later timestamps are likely caused by slightly overestimated wind speeds in the ERA5 data and the inability to account for small-scale local wind variations.

This analysis demonstrates that contrail tracking can provide valuable information to further narrow down the number of potential contrail sources, particularly by identifying well-aligned matches at the first contrail

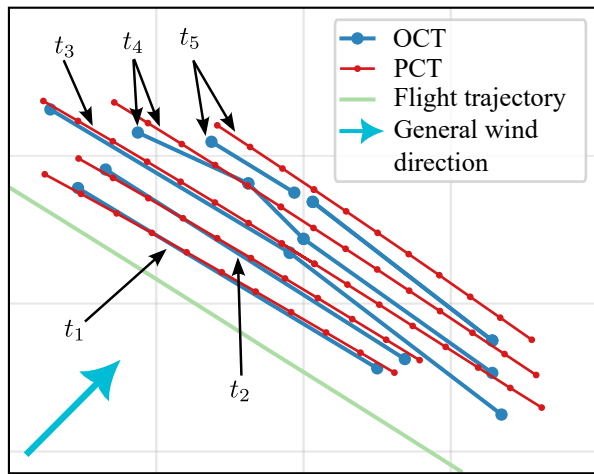


FIG 13. Comparison of the OCT and the PCT of the most probable candidate flight from the case study across multiple timestamps.

observation, ruling out emissions that occurred after the first observation, and excluding flights that only matched isolated segments. However, it also shows that relying on a single timestamp is often insufficient for accurately determining contrail sources, as PCT–OCT deviations can increase with the time between formation and observation, sometimes making them indistinguishable from falsely matched flights.

4. CONCLUSION AND OUTLOOK

Conclusion

This study developed an automated method to associate contrails detected in geostationary satellite imagery with their source flights by applying a two-stage filtering process, a geographic scoring metric, and a backtracking procedure. The results show improved performance compared to an existing method, producing fewer but more plausible matches, mainly due to the implementation of the second filtering stage, which calculates wind displacement using ERA5 meteorological data.

A statistical evaluation demonstrated that the method can reliably generate a small set of likely candidate flights. Analysis of the environmental conditions at formation further suggests that real originators are indeed included within these matches. Moreover, higher match quality scores were shown to correlate with a greater probability of true matches, underlining the usefulness of this metric. However, about 30% of contrails could not be matched to any flight, which may be due to false detections or to inaccuracies in wind data that accumulate over time. Addressing these issues and improving contrail detection models represents a sensible next step.

A case study of the same contrail at different times showed that tracking a contrail across multiple images

provides additional information that can be used for further filtering. Automating this temporal tracking and developing a dedicated filtering logic would offer significant further improvements in accuracy and reliability.

Outlook

Overall, the current approach is not yet capable of creating a large database of real-world observations for the validation of contrail prediction and simulation models. However, the method provides a foundation for future development, and the results indicate that with further improvements this goal may be achievable. The first step is to automate the contrail-tracking methodology. Both the presented example and the case reported in the literature demonstrate that this approach can further narrow down the number of candidate flights per contrail. A second priority is to enhance the contrail detection model to cover additional geographic regions. Although this expansion might require retraining due to different satellite sensors technology, it will also increase the flexibility of the method and enlarge the underlying data base. Regarding the evaluation component, enhancements to the SAC metric are feasible if the presently employed flight trajectories are linked to a comprehensive mission-level analysis. Such an analysis can provide a thrust-requirement depending on aircraft type and flight length, which can then be used in a thermodynamic cycle model to produce more realistic figures—e.g., the engine’s thermal efficiency.

Contact address:

simon.kroeger@tu-braunschweig.de

Acknowledgements

We would like to thank the Förderverein Strömungsmaschinen C. Pfeleiderer e. V. for their funding.

References

- [1] Ernst Schmidt. Die Entstehung von Eisnebel aus den Auspuffgasen von Flugmotoren, 1941. Issue: 44 Pages: 1–15 Series: Schriften der Deutschen Akademie der Luftfahrtforschung Volume: 5. <https://elib.dlr.de/107948/>.
- [2] H. Appleman. The Formation of Exhaust Condensation Trails by Jet Aircraft. *Bulletin of the American Meteorological Society*, 34(1):14–20, Jan. 1953. ISSN: 0003-0007, 1520-0477. DOI: [10.1175/1520-0477-34.1.14](https://doi.org/10.1175/1520-0477-34.1.14).
- [3] Ulrich Schumann. On conditions for contrail formation from aircraft exhausts. *Meteorologische Zeitschrift*, 5:4–23, 1996. Publisher: Borntraeger Science Publishers.
- [4] Guy P. Brasseur, Mohan Gupta, Bruce E. Anderson, Sathya Balasubramanian, Steven Bar-

- rett, David Duda, Gregg Fleming, Piers M. Forster, Jan Fuglestedt, Andrew Gettelman, Rangasayi N. Halthore, S. Daniel Jacob, Mark Z. Jacobson, Arezoo Khodayari, Kuo-Nan Liou, Marianne T. Lund, Richard C. Miake-Lye, Patrick Minnis, Seth Olsen, Joyce E. Penner, Ronald Prinn, Ulrich Schumann, Henry B. Selkirk, Andrei Sokolov, Nadine Unger, Philip Wolfe, Hsi-Wu Wong, Donald W. Wuebbles, Bingqi Yi, Ping Yang, and Cheng Zhou. Impact of Aviation on Climate: FAA's Aviation Climate Change Research Initiative (ACCR) Phase II. *Bulletin of the American Meteorological Society*, 97(4):561–583, Apr. 2016. ISSN: 0003-0007, 1520-0477. DOI: [10.1175/BAMS-D-13-00089.1](https://doi.org/10.1175/BAMS-D-13-00089.1).
- [5] D.S. Lee, D.W. Fahey, A. Skowron, M.R. Allen, U. Burkhardt, Q. Chen, S.J. Doherty, S. Freeman, P.M. Forster, J. Fuglestedt, A. Gettelman, R.R. De León, L.L. Lim, M.T. Lund, R.J. Millar, B. Owen, J.E. Penner, G. Pitari, M.J. Prather, R. Sausen, and L.J. Wilcox. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244:117834, Jan. 2021. ISSN: 13522310. DOI: [10.1016/j.atmosenv.2020.117834](https://doi.org/10.1016/j.atmosenv.2020.117834).
- [6] Intergovernmental Panel On Climate Change (Ipcc). *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 1 edition, July 2023. ISBN: 978-1-009-15789-6. DOI: [10.1017/9781009157896](https://doi.org/10.1017/9781009157896).
- [7] Klaus Gierens, Ling Lim, and Kostas Eleftheratos. A Review of Various Strategies for Contrail Avoidance. *The Open Atmospheric Science Journal*, 2(1):1–7, Feb. 2008. ISSN: 18742823. DOI: [10.2174/1874282300802010001](https://doi.org/10.2174/1874282300802010001).
- [8] Esther Roosenbrand, Junzi Sun, and Jacco Hoekstra. Contrail minimization through altitude diversions: A feasibility study leveraging global data. *Transportation Research Interdisciplinary Perspectives*, 22:100953, Nov. 2023. ISSN: 25901982. DOI: [10.1016/j.trip.2023.100953](https://doi.org/10.1016/j.trip.2023.100953).
- [9] U. Schumann. A contrail cirrus prediction model. *Geoscientific Model Development*, 5(3):543–580, May 2012. ISSN: 1991-9603. DOI: [10.5194/gmd-5-543-2012](https://doi.org/10.5194/gmd-5-543-2012).
- [10] Joe Yue-Hei Ng, Kevin McCloskey, Jian Cui, Vincent R. Meijer, Erica Brand, Aaron Sarna, Nita Goyal, Christopher Van Arsdale, and Scott Geraedts. OpenContrails: Benchmarking Contrail Detection on GOES-16 ABI, 2023. Version Number: 2. DOI: [10.48550/ARXIV.2304.02122](https://doi.org/10.48550/ARXIV.2304.02122), <https://arxiv.org/abs/2304.02122>.
- [11] Emmanuel Riggi-Carolo, Thomas Dubot, Claire Sarrat, and Judicaël Bedouet. AI-Driven Identification of Contrail Sources: Integrating Satellite Observation and Air Traffic Data. *Journal of Open Aviation Science*, 1(2), Oct. 2023. ISSN: 2773-1626. DOI: [10.59490/joas.2023.7209](https://doi.org/10.59490/joas.2023.7209).
- [12] Rémi Chevallier, Marc Shapiro, Zebediah Engberg, Manuel Soler, and Daniel Delahaye. Linear Contrails Detection, Tracking and Matching with Aircraft Using Geostationary Satellite and Air Traffic Data. *Aerospace*, 10(7):578, June 2023. ISSN: 2226-4310. DOI: [10.3390/aerospace10070578](https://doi.org/10.3390/aerospace10070578).
- [13] Matthias Schafer, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. Bringing up OpenSky: A large-scale ADS-B sensor network for research. In *IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*, pages 83–94, Berlin, Germany, Apr. 2014. IEEE. ISBN: 978-1-4799-3147-7 978-1-4799-3146-0. DOI: [10.1109/IPSN.2014.6846743](https://doi.org/10.1109/IPSN.2014.6846743).
- [14] Xavier Olive. traffic, a toolbox for processing and analysing air traffic data. *Journal of Open Source Software*, 4(39):1518, July 2019. ISSN: 2475-9066. DOI: [10.21105/joss.01518](https://doi.org/10.21105/joss.01518).
- [15] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnoti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sebastien Villaume, and Jean-Noël Thépaut. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, July 2020. ISSN: 0035-9009, 1477-870X. DOI: [10.1002/qj.3803](https://doi.org/10.1002/qj.3803).
- [16] Roger Teoh, Ulrich Schumann, Edward Gryspeerdt, Marc Shapiro, Jarlath Molloy, George Koudis, Christiane Voigt, and Marc E. J. Stettler. Aviation contrail climate effects in the North Atlantic from 2016 to 2021. *Atmospheric Chemistry and Physics*, 22(16):10919–10935, Aug. 2022. ISSN: 1680-7324. DOI: [10.5194/acp-22-10919-2022](https://doi.org/10.5194/acp-22-10919-2022).
- [17] Roger Teoh, Zebediah Engberg, Ulrich Schumann, Christiane Voigt, Marc Shapiro, Susanne Rohs, and Marc E. J. Stettler. Global aviation contrail climate effects from 2019 to 2021. *Atmospheric Chemistry and Physics*,

24(10):6071–6093, May 2024. ISSN:1680-7324.
DOI: [10.5194/acp-24-6071-2024](https://doi.org/10.5194/acp-24-6071-2024).

- [18] Marc Shapiro, Zeb Engberg, Roger Teoh, Marc Stettler, and Tom Dean. pycontrails: Python library for modeling aviation climate impacts. *Zenodo [code]*, 10, 2023.
- [19] Junzi Sun. Neural network models for contrail detection and segmentation. Available online at: <https://github.com/junzis/contrail-seg>, last accessed on 31.01.2025. <https://github.com/junzis/contrail-seg>.
- [20] Yan Xiaolong. Skeleton Network. Available online at: <https://github.com/Image-Py/sknw>, last accessed on 31.01.2025.
- [21] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, 1 edition, Aug. 2002. ISBN:978-0-521-81087-6 978-0-521-00924-9 978-0-511-79125-3. DOI: [10.1017/cbo9780511791253](https://doi.org/10.1017/cbo9780511791253).