



GVCCS: A Dataset for Contrail Identification and Tracking on Visible Whole Sky Camera Sequences

Gabriel Jarry¹, Ramon Dalmau¹, Philippe Very¹, Franck Ballerini¹, and Stefania-Denisa Bocu¹

¹EUROCONTROL, Aviation Sustainability Unit (ASU), Aerodrome Centre Bois des Bordes, 91220 Brétigny-sur-Orge, France

Correspondence: Gabriel Jarry (gabriel.jarry@eurocontrol.int)

Abstract. Aviation's climate impact includes not only CO₂ emissions but also significant non-CO₂ effects, especially from contrails. These ice clouds can alter Earth's radiative balance, potentially rivalling the warming effect of aviation CO₂. Physics-based models provide useful estimates of contrail formation and climate impact, but their accuracy depends heavily on the quality of atmospheric input data and on assumptions used to represent complex processes like ice particle formation and humidity-driven persistence. Observational data from remote sensors, such as satellites and ground cameras, could be used to validate and calibrate these models. However, existing datasets don't explore all aspect of contrail dynamics and formation: they typically lack temporal tracking, and do not attribute contrails to their source flights. To address these limitations, we present the Ground Visible Camera Contrail Sequences, a new open data set of contrails recorded with a ground-based all-sky camera in the visible range. Each contrail is individually labelled and tracked over time, allowing a detailed analysis of its lifecycle. The dataset contains 122 video sequences (24,228 frames) and includes flight identifiers for contrails that form above the camera. As reference, we also propose a unified deep learning framework for contrail analysis using a panoptic segmentation model that performs semantic segmentation (contrail pixel identification), instance segmentation (individual contrail separation), and temporal tracking in a single architecture. By providing high-quality, temporally resolved annotations and a benchmark for model evaluation, our work supports improved contrail monitoring and will facilitate better calibration of physical models. This sets the groundwork for more accurate climate impact understanding and assessments.

1 Introduction

Aviation contributes to global climate change not only through carbon dioxide (CO₂) emissions but also through a variety of non-CO₂ effects, including nitrogen oxides (NO_x), water vapour and aerosols. Among these, condensation trails (contrails), ice-crystal clouds formed by aircraft at typical cruising altitudes, stand out for their potentially large, yet uncertain, radiative impact. Though they often appear as ephemeral white streaks in the sky, persistent contrails can spread into extensive cirrus-like cloud formations that trap outgoing long-wave radiation, warming the planet. Recent studies suggest that the climate forcing due to contrail cirrus clouds is of the same order of magnitude as aviation CO₂ emissions (Lee et al., 2021; Teoh et al., 2023), although this depends on the metric used (Borella et al., 2024).



Yet, accurately assessing the climate impact of contrails remains a significant challenge for both aviation and climate sci-
25 entists. The lifecycle of the contrails depends on complex interrelated processes, such as ice nucleation, crystal growth, wind-
driven dispersion, and interaction with natural clouds, that are sensitive to ambient atmospheric conditions. Small variations
in temperature and humidity, particularly relative humidity with respect to ice, can determine whether a contrail dissipates
quickly or persists and spreads. This sensitivity, combined with the diurnal variability in radiative forcing (cooling when re-
flecting sunlight during the day; warming when trapping infrared radiation at night), makes the net effect of contrails both
30 context-dependent and extremely difficult to model reliably.

While contrail impacts have traditionally been studied using physical models, recent advances in remote sensing and com-
puter vision now offer a valuable observational perspective. Physics-based models, such as the Contrail Cirrus Prediction model
(CoCiP) (Schumann, 2012) or Aviation Particle Composition and Evolution Model for Microphysics (APCEMM) (Fritz et al.,
2020) simulate the lifecycle of a contrail by solving complex equations that describe the interaction between aircraft emissions
35 and atmospheric conditions. These models provide valuable theoretical insights, but their accuracy is heavily dependent on
the quality of the input data (Gierens et al., 2020). Key parameters, such as atmospheric temperature, humidity, and aircraft
engine characteristics, are often uncertain and these uncertainties propagate through the calculations, affecting the reliability of
the results. Moreover, detailed simulations of contrail microphysics and radiative effects can be computationally demanding,
particularly when applied to global-scale analyses.

40 Observational methods, using satellite and ground-based imagery, offer a direct and data-driven way to study contrails, com-
plementing theoretical models. Advances in high-resolution remote sensing and computer vision have made these methods
increasingly effective (Meijer et al., 2022; McCloskey et al., 2021; Ng et al., 2023; Chevallier et al., 2023). Beyond detec-
tion, observational data should play an increasing role in the future in refining physics-based models by providing empirical
validation and calibrating the uncertain parameters mentioned above.

45 Integrating observational data with air traffic information like Automatic Dependent Surveillance-Broadcast (ADS-B) and
meteorological data holds significant promise for advancing our understanding of the contrail lifecycle and climate impact.
Linking contrails to specific flights, for which detailed parameters (e.g., engine type, altitude, and atmospheric conditions) are
known, will allow for a better understanding of the role of these parameters into contrail formation and dynamics. However,
achieving this integration requires addressing foundational challenges: accurately identifying contrails in images and/or videos,
50 distinguishing them from natural clouds (semantic segmentation), detecting individual instances (instance segmentation), and
tracking their evolution over time. This paper focuses on these critical first steps, developing robust methods for contrail
segmentation and tracking in both individual ground camera images and videos. While attribution remains a very challenging
task to perform at scale using in particular geostationary satellites (Chevallier et al., 2023; Riggi-Carolo et al., 2023; Geraedts
et al., 2024; Sarna et al., 2025) our work provide the necessary tools to reliably detect and track contrails locally, laying the
55 groundwork for subsequent integration with flight and meteorological data.

Despite growing interest in observational contrail analysis, publicly available datasets are still limited in scope. The most
universally used resource, Google's OpenContrails, offers instance-level masks only on the central GOES-16 frame, with
surrounding images left unannotated, hindering contrail tracking across time. In contrast, Sarna et al. (2025) introduced Syn-



thOpenContrails, which overlays synthetic contrails and annotations onto real scenes, providing full per-frame localization, tracking, and flight attribution, demonstrating that richly annotated data exists, even if confined to synthetic contrail overlays rather than human annotation. An ideal scenario would be a fully annotated video dataset where every frame is humanly labelled and each contrail is assigned a persistent identifier.

To advance research in this area, this paper presents the Ground Visible Camera Contrail Sequences (GVCCS), an open dataset (Jarry et al., 2025) with instance-level annotations, derived from ground-based video recordings in Brétigny-sur-Orge, France (Réuniwatt CamVision visible ground-based camera). Our dataset includes 122 videos (of duration between 20 min to five hours) with a total frame number of around 24,200, each annotated with instance-level labels. By making this dataset openly available, this paper provides a valuable benchmark for both the atmospheric and aviation research communities.

To support future performance comparisons, we introduce here a deep learning-based model for contrail segmentation and tracking. Instead of relying on separate models for these tasks, an approach that often requires complex, ad-hoc combinations of techniques, we adopt a unified framework based on Mask2Former, a state-of-the-art computer vision model. Mask2Former is designed for panoptic segmentation, which combines semantic segmentation (labelling each pixel with a class, e.g., *contrail* or *sky*) and instance segmentation (distinguishing between individual objects, e.g., different contrails). In addition to separating contrails from clear sky, it could handle complex backgrounds, such as low-altitude cloud layers that partially or fully obscure contrails, by assigning appropriate *cloud* labels while still maintaining unique instance identities. For example, in a single image, panoptic segmentation can identify all visible contrail pixels, correctly label intervening clouds, and assign consistent instance masks to each contrail, even when they overlap, intersect, appear fragmented, or are seen through thin cloud cover. In fact, contrails often break into multiple disconnected components due to atmospheric conditions and natural dissipation processes. A robust monitoring system must not only identify these fragments, but also associate them with the correct contrail instance.

It is worth noting that, fragmentation poses a significant challenge for contrail analysis based solely on images or videos: visually disjointed segments from the same flight must be grouped without external data. Moreover, low-altitude cloud obscuration and sun glare can further interrupt or mask contrail continuity, producing multi-polygon annotations even for a single physical contrail. In operational settings, however, it is possible to first perform single-polygon instance segmentation and then associate multiple instances with the same flight using auxiliary data such as aircraft trajectories and wind fields. This post-processing step enables grouping across time and space based on flight identity rather than visual continuity. In this work, we restrict ourselves to purely image-based analysis and defer the integration of external data sources to future work.

Mask2Former, originally designed for individual images, can be easily extended to video data to improve the consistency of panoptic segmentation across frames. By leveraging temporal information, Mask2Former for videos performs semantic segmentation, instance segmentation, and tracking in an integrated manner. In this paper, we study both the frame-based and video-based versions of Mask2Former, comparing their performance on our dataset.

The remainder of this paper is structured as follows. Section 2 provides the necessary background on contrail formation and computer vision techniques, establishing the foundation for the challenges addressed in this work. Section 3 reviews related work on contrail datasets and segmentation models, highlighting current limitations and motivating our approach. Section 4



introduces our newly developed video-based dataset, detailing its annotation methodology and unique instance-level structure.

95 Section 5 describes our panoptic segmentation framework based on the Mask2Former architecture. Section 6 presents and analyses the experimental results. Finally, Section 7 summarises our main contributions and outlines future research directions.

2 Background

This section introduces the key concepts necessary to understand the challenges addressed in this work. We begin by outlining the physical processes behind contrail formation and the implications for climate, focusing on why contrails are particularly
100 difficult to detect and track. We then review relevant computer vision techniques, specifically object detection and image segmentation, and assess their suitability for analysing contrails.

2.1 The science of contrails

Contrails are artificial clouds that form behind aircraft when hot, humid engine exhaust mixes with the cold, low-pressure air at cruising altitudes, typically in the range 8-12 km. If the atmospheric conditions are right, specifically, if the temperature
105 falls below -40°C and the air is sufficiently humid, the water vapour in the exhaust condenses and freezes into ice crystals. This process, modelled and quantified by the Schmidt-Appleman criterion (Appleman, 1953), produces the familiar thin, white trails visible in the sky. Some contrails dissipate rapidly, while others persist and spread, eventually forming larger ice cloud structures known as contrail cirrus.

Like natural clouds, contrails influence the Earth's radiation budget: they trap outgoing long-wave radiation, leading to warm-
110 ing, while also reflecting incoming solar radiation, which has a cooling effect. The net result depends on the contrail's altitude, optical properties, lifespan, and the time of day. The precise relative impact depends on the climate metric chosen (Borella et al., 2024); however, contrails are thought to warm the climate at a level of the same order of magnitude as aviation's CO_2 emissions (Lee et al., 2021; Teoh et al., 2023). This makes the monitoring and characterization of contrails an essential part of understanding aviation's full environmental footprint (Teoh et al., 2023) and developing mitigation strategies (Teoh et al.,
115 2020).

As mentioned above, the observational viewpoint offers an alternative perspective that focuses on detecting and analysing contrails directly in atmospheric imagery. However, detecting and tracking contrails presents several technical challenges, which helps explain the growing research interest in the topic. Satellite imagery often lacks the spatial and temporal resolution needed to detect contrails in their early stages (Ng et al., 2023). Geostationary satellites have a nominal spatial resolution of
120 about 0.5-2 km and a temporal resolution of 5-15 min, which is often insufficient to capture the narrow, faint, and short-lived nature of freshly formed contrails unless they persist and grow. Even when contrails do spread into detectable cloud structures, they are difficult to distinguish from natural cirrus, particularly in scenes with complex cloud layers. Moreover, by the time a contrail is visible in satellite images, it has often drifted and deformed, complicating the attribution to the flight that produced it (Chevallier et al., 2023; Sarna et al., 2025). This linkage is crucial, as identifying the originating flight enables researchers to



125 retrieve essential details such as aircraft type and engine model, key inputs for assessing contrails' environmental impact and improving physical models through comparison with empirical observations.

Ground-based cameras (Schumann et al., 2013; Low et al., 2025) offer a complementary perspective with critical advantages. Positioned beneath flight paths, these systems can capture high-resolution images and video with far greater spatial and temporal fidelity than satellites. Crucially, they can detect contrails immediately after formation, while they are still thin, linear, and visually distinct. This early visibility simplifies the task of associating observed contrails with the specific flight responsible, especially when combined with precise trajectory data. The main drawback is, naturally, their restricted coverage, which hinders the ability to monitor contrails from their formation to dissipation.

While not the focus of this paper, one promising direction involves combining ground-based and satellite observations into a unified monitoring framework. In such a system, contrails would first be detected in high-resolution ground-based imagery and attributed to specific flights using trajectory and weather data providing access to key aircraft and engine parameters. Crucially, to enable continuous tracking beyond the limited field of view of the ground-based camera, these contrails would then need to be reliably linked to their evolving counterparts in satellite imagery as they drift, expand, and age. Successfully associating contrails across these two modalities, ground and satellite, would allow monitoring of their full lifecycle from formation to dissipation while preserving information about the specific aircraft and flight responsible for creating them.

140 2.2 Computer vision techniques for contrail monitoring

Contrails are visually challenging targets for computer vision due to their thin, elongated shapes, variable curvature, and tendency to fragment or fade over time. These characteristics make them fundamentally different from the objects typically addressed in standard object detection benchmarks, such as vehicles and animals in datasets like the Common Objects in Context (COCO) dataset, which features well-defined, discrete objects.

145 Traditionally, object detection methods localise targets using bounding boxes, usually axis-aligned rectangles. This approach works well for objects like cars or animals, which are compact and roughly rectangular, but performs poorly for contrails. A single axis-aligned bounding box may inadvertently include multiple contrail segments or large amounts of background sky, while missing parts of curved or fragmented trails. Oriented bounding boxes offer some improvement by allowing rotation, which better fits the geometry of elongated contrails. However, they still fall short in capturing fine-grained shapes, gaps, or fading segments. Figure 1 shows the limitations of axis-aligned and oriented bounding boxes for object detection on contrails.

Instance segmentation provides a more precise solution by predicting pixel-level masks for each individual object. This approach is particularly beneficial for contrails, as it can delineate each trail accurately even when they intersect, overlap, or dissipate unevenly. For instance, two overlapping contrails that fade at different rates can still be assigned to distinct instances.

Semantic segmentation, in contrast, labels each pixel by class, e.g., *contrail* or *sky*, but does not distinguish between individual contrails. This is insufficient when studying temporal evolution or interactions between specific contrails, since it treats all contrails as a single undifferentiated class.

Panoptic segmentation combines the strengths of both approaches: it assigns a class label to every pixel (semantic segmentation) and an instance identifier where appropriate (instance segmentation). In this framework, things such as individual




Figure 1. Illustration of bounding box detection on contrails. Each detected contrail is highlighted with a distinct colour. It should be noted that elongated or fragmented contrails challenge bounding box alignment and separation.

contrails are assigned unique instance labels, while stuff like the background sky or natural clouds is labelled only by class. This unified view is well-suited to contrail monitoring, enabling fine-grained analysis of individual contrails within the broader atmospheric context. Moreover, the framework can be readily extended to additional classes (e.g., cirrus, cumulus) for more comprehensive scene understanding, provided, of course, that these classes have been effectively and consistently labelled during dataset creation, which introduces an additional layer of complexity to the annotation campaign. Figure 2 illustrates the instance, semantic and panoptic segmentation methods.




Figure 2. Comparison of segmentation methods applied to illustrative contrails. Distinct colours indicate different contrail instances or classes, depending on the method.



165 An important but often overlooked issue in the literature is how contrails are geometrically represented. In reality, a single contrail may consist of several disconnected segments, for example, due to fading or occlusion, making it a multi-polygon shape. See, for instance, the green contrail in Figs. 2a and 2c. However, the most natural approach is to simplify this by treating each segment as a separate, independent polygon, effectively assuming that each fragment belongs to a different contrail.

While this simplification avoids the complexity of handling multi-polygons directly, it introduces a significant challenge: to
170 reconstruct the full contrail, one must find a way to link fragmented pieces together. This requires ad-hoc linking strategies, which vary in complexity and accuracy. Some methods rely purely on the geometric properties of the fragments, such as their proximity or alignment, while others incorporate external data, such as aircraft flight paths or meteorological information, to make more informed associations.

In this work, we adopt panoptic segmentation as the foundation for segmenting and tracking contrails. This choice is mo-
175 tivated by its ability to simultaneously achieve instance-level precision and maintain contextual awareness of the surrounding scene. Moreover, by explicitly addressing the issue of fragmented contrails, our method enables instance-level identification of contrails without requiring external sources of information, such as flight or weather data. This is particularly valuable in scenarios where such data may be unavailable or incomplete. However, we also explore an alternative version of the model that treats each contrail fragment as an independent instance, under the assumption that a downstream algorithm, leveraging
180 external traffic and meteorological data, will later associate these segments with their corresponding flights. The comparative evaluation of these two strategies, self-contained instance identification versus externally supported post-association, will be presented in future publications. In this paper, we focus solely on presenting the contrail segmentation models themselves.

3 State of the art

This section presents an overview of prior work in contrail segmentation and analysis, focusing first on the datasets that have
185 been developed to support this research, and then on the computational models used for contrail segmentation and flight attribution. The scope and key features of existing datasets are outlined, with particular attention given to the limited availability of temporal annotations and flight attribution ground truth. Subsequently, we examine state-of-the-art segmentation and tracking methods, particularly deep learning-based approaches, assessing their applicability and performance in contrail analysis. This review highlights gaps in current research and motivates the contributions presented in this paper.

190 3.1 Datasets

Recent advances in contrail detection have been supported by the development of annotated datasets, primarily based on satellite imagery. These datasets have facilitated the application of computer vision techniques for contrail identification, although aspects such as temporal continuity and integration with flight metadata remain limited in most cases. In this section, we review the most relevant publicly available datasets and place our contributions within this context.

195 Kulik (2019) and Meijer et al. (2022) are to our knowledge the first studies to leverage a modern, data-driven, deep learning framework for large-scale contrail segmentation. The authors developed and applied convolutional neural networks, which



were trained using a manually curated dataset comprising over 100 manually annotated geostationary GOES satellite images with instance segmentation.

One of the first large-scale labelling efforts in contrail detection was led by Google Research, beginning with the development of a contrail dataset based on high-resolution Sentinel satellite imagery (McCloskey et al., 2021). Human experts manually annotated the images using structured guidelines, producing polygonal masks for each visible contrail segment. Multiple annotators independently labelled each image, and the dataset includes all individual annotations, with the option to filter results by majority consensus. This methodology improved both the spatial precision and overall quality of the labels.

Building on this work, Google released the OpenContrails dataset (Ng et al., 2023), which is based on images from the GOES-16 Advanced Baseline Imager (ABI). Thanks to the 10-minute temporal resolution provided by the geostationary orbit of GOES-16, the dataset is well suited to study contrails at large scales. OpenContrails offers temporal context by including short sequences of unlabelled images surrounding each annotated frame, providing valuable information to annotators for more accurate labelling. Only the central frame in each sequence is annotated, therefore not allowing a direct comparison of contrail dynamics with physical models. Notably, a 2025 update introduced instance-level labels, enabling the use of the dataset for instance-based models and expanding its potential for more advanced contrail analysis.

Gourgue et al. (2025) introduce an open-access corpus of around 1,600 polygon-annotated hemispheric sky images acquired at the SIRTa atmospheric laboratory, near Paris, offering class labels that distinguish *young*, *old*, and *very old* contrails as well as several confounding artefacts. By capturing high-resolution ground views minutes after formation, the dataset fills the temporal-spatial gap left by satellite benchmarks.

Rather than creating a dataset for training modern convolutional networks on segmentation tasks, Low et al. (2025) manually annotated the correspondence between contrail waypoints, derived from the application of the CoCiP model and observations from their wide-angle ground camera system. This approach is particularly well-suited for directly assessing and parametrizing physical models.

Meijer et al. (2024) is to our knowledge the first example of dataset collocating images on two different remote sensors: they assembled a dataset specifically for contrail altitude-altitude estimation, comprising over 3,000 cases over the contiguous United States (2018-2022). Contrails were first located via automated detection in GOES-16 ABI infrared imagery, then precisely collocated, correcting for parallax and wind advection, with CALIOP lidar cross-sections. The team then conducted manual inspections of the matched imagery to verify and validate alignment. This benchmark dataset linking geostationary contrail signatures to high-resolution vertical profiles enables supervised deep-learning approaches to predict contrail top heights from ABI data.

A significant advance in contrail detection has been the development of synthetically labelled datasets. Chevallier et al. (2023) generated a synthetic dataset using CoCiP (Schumann, 2012) to overlay contrail polygons onto GOES-16 imagery, enabling the first instance segmentation pipeline for contrail detection. The performance of flight assignment algorithms was validated using actual GOES data, through manual inspection rather than synthetic reference ground truth. Building on this synthetic foundation, Sarna et al. (2025) introduced a benchmark dataset, SynthOpenContrails, with sequences of synthetic contrail detections tied to known flight metadata, providing the first opportunity to quantitatively evaluate and improve con-



trail-flight attribution algorithms. To our knowledge, this is the only dataset providing localized and tracked contrails with attributable ground truth, albeit synthetic. While the use of synthetic datasets represents a modern and cutting-edge technique for training algorithms, the use of manually labelled data as test sets is still theoretically preferable to objectively assess algorithmic performance. However, obtaining such datasets on geostationary satellite images, with their coarse resolution, remains very difficult at this stage, which motivates the approach adopted by the authors. As mentioned by Sarna et al. (2025), obtaining such a reference dataset with ground truth for flight attribution based on human annotations is definitely feasible in principle with higher resolution low orbit satellites or ground-based cameras, which is the focus of the present work.

Overall, while existing datasets have contributed valuable resources there is a lack of comprehensive, human-labelled data containing temporally resolved, instance-level and flight-attributed annotations. Our work addresses this issue by introducing a dataset designed to provide these annotations, collected using our ground camera system.

3.2 Models

Contrail monitoring with computer vision was first pioneered in the early nineties (Forkert et al., 1993; Mannstein et al., 1999), using non-data-driven image-analysis techniques. Their work applied linear-kernel methods, direct thresholding of brightness temperature difference channels, and early Hough-transform operators (Pratt, 2007) optimized for linear shape detection, to identify contrails in AVHRR satellite imagery. This approach was further improved by Vazquez-Navarro et al. (2010) and Duda et al. (2013).

To the best of our knowledge, Kulik (2019); Meijer et al. (2022) represent the earliest applications of modern convolutional networks to pixel-level classification and semantic segmentation. Building on the OpenContrails dataset, Ng et al. (2023) employed semantic segmentation algorithms, specifically DeepLabV3 (Chen et al., 2017, 2018), to identify contrails in ash-rgb composites using brightness temperature differences. Their work demonstrated that adding temporal context via a 3D encoder, incorporating the time dimension, led to improved performance. Moreover, results from the subsequent Kaggle competition showed that UNet models (Ronneberger et al., 2015) equipped with modern transformer backbones, such as MaxViT (Tu et al., 2022) and CoatNet (Dai et al., 2021), achieved even stronger results (Jarry et al.).

Using an ensemble approach, Ortiz et al. (2025) combined six neural networks, including U-Net, DeepLab, and transformer architectures, and applied optical-flow-based corrections to maintain temporal consistency across consecutive satellite frames. Meanwhile, Sun and Roosenbrand (2025) introduced a Hough-space line-aware loss for few-shot scenarios, supplementing Dice loss with a global alignment term to encourage predictions to align with linear structures.

Shifting from pixel-level masks to instance-level contrail segmentation and making use of synthetic data Chevallier et al. (2023) introduced the first algorithmic pipeline focused on instance segmentation for contrail detection, utilizing the Mask R-CNN algorithm (He et al., 2017). Similarly, Van Huffel et al. (2025) adopted Mask R-CNN to process images captured by their wide-angle ground camera system.

The difficult task of attributing detected contrails to individual flights (typically using ADS-B information) in geostationary satellite imagery has been the focus of several recent studies. Chevallier et al. (2023) introduced a pipeline that combines contrail detection, tracking, and matching with aircraft using geometric criteria and wind-corrected trajectories. Riggi-Carolo



et al. (2023) proposed an probabilistic matching methods that account for uncertainties in flight data and atmospheric conditions leveraging as well on Hough-based line detection. Geraedts et al. (2024) presented a scalable system designed to assign contrails to flights on a large scale, enabling routine monitoring of contrail formation and supporting climate assessments. Sarna et al. (2025) systematically benchmarked and refined these attribution algorithms, highlighting common challenges and proposing improved association metrics, building on the release of the synthetically generated SynthOpenContrails dataset

By contrast, our work targets ground-based imagery, capturing contrails immediately after formation and enabling near-instantaneous flight attribution via ADS-B data. We harness panoptic segmentation using Mask2Former, trained on high-resolution video, to extract pixel-accurate masks of individual contrails and track them over time. This fills the gap in early-stage contrail detection and provides richer spatial and temporal detail than existing satellite-based models.

275 4 Dataset

The primary contribution of this paper is the introduction of a new dataset designed to support contrail detection, tracking and attribution. This section provides a detailed overview of the dataset. Section 4.1 describes the data collection and labelling campaign. Section 4.2 summarizes the structure and content of the dataset.


4.1 Data collection and labelling campaign

280 To support the development of machine learning models for contrail detection, we conducted an extensive labelling campaign as part of the ContrailNet project. Visible-spectrum image sequences were acquired using a all-sky ground-based camera installed on the roof of the EUROCONTROL Innovation Hub, capturing the sky every 30 s at a resolution of $1,976 \times 2,032$ pixels.

Our camera provider, Reuniwatt, has delivered a dual all-sky camera system: the first unit, CamVision, operates in the visible spectrum, capturing high-resolution fisheye images every 30 s with on-board processing and self-calibration, ensuring reliable daytime operation even in dusty or wet conditions. The second unit, SkyInsight, uses long-wave infrared (8-13 μm) imaging via a chrome-coated hemispherical mirror and will be used in future research.

The raw all-sky images were first geometrically projected onto a square grid. This projection process, uses camera-specific calibration files to associate each pixel with its corresponding azimuth and zenith angles, effectively removing lens distortions and re-mapping the sky onto a uniform Cartesian representation. A 75×75 km grid of georeferenced points was computed at a fixed cloud altitude (10 km), and a linear interpolation scheme was used to assign raw pixel values to the projected frame. The output is a square image of size $1,024 \times 1,024$ pixels that preserves the spatial geometry of the sky above the camera.

To improve the visual clarity and consistency of the sequences, each projected image then undergoes a three-step enhancement process. First, brightness is increased using a linear scaling operation. Second, local contrast is enhanced via CLAHE (Contrast Limited Adaptive Histogram Equalization), which boosts features like contrails without overexposing the image. Finally, colour warmth is reduced by rebalancing the blue and red channels, improving contrail visibility in high-sunlight con-



(a) Raw image from the ground-based camera with visible contrails. (b) Geometrically projected and enhanced image used for annotation.

Figure 3. Side-by-side comparison of the raw ground-based camera image and the geometrically projected version used for annotation. Projection corrects perspective distortions and enhances contrast for better segmentation.

ditions. This preprocessing pipeline proved essential in highlighting fine contrail structures, especially in complex atmospheric scenes. Both raw and enhanced projected images are displayed in Fig. 3.

The process of labelling was applied to video sequences, each sequence comprised between 60 and 480 images, corresponding to durations of 30 min to 4 h, enabling the temporal tracking of contrails throughout their formation and dissipation phases.

The labelling process was carried out using a dedicated annotation tool developed by Encord, who also provided a professional team of annotators. We maintained close collaboration with this team through regular coordination meetings, during which the annotation guide was developed and iteratively refined. The labelling platform was specifically configured to overlay flight trajectory data above the camera's field of view, assisting annotators in identifying *new* contrails, those forming above the camera and visibly associated with a known aircraft trajectory. In contrast, *old* contrails were defined as those already present at the start of a sequence or likely formed outside the camera's field of view, making flight association impossible.

Each contrail was annotated using high-precision polygons that tracked its spatial extent throughout its visible evolution, from early linear stages to advanced spreading phases. When contrails became fragmented or partially obstructed by clouds, multiple polygons were used and linked using relational attributes (*fragmented contrail* and *cloud obstruction*) to preserve temporal continuity.

To ensure the highest annotation quality, the campaign incorporated a multi-stage review protocol. An initial calibration phase was conducted using a sample dataset to harmonise interpretation and identify edge cases. Each labelled sequence then



Table 1. Descriptive statistics of the annotated contrail dataset

Metric	Value
Total sequences (labelled)	122
Total images	24,228
Average sequence duration	96.6 min
Images per sequence (minimum / maximum / mean)	41 / 600 / 198.6
Total annotated contrail instances	4,651
Total unique flight IDs assigned	3,354
Total polygons annotated	176,234
Contrail duration (minimum / maximum / mean)	0.5 / 142.5 / 14.6 min
Polygons per contrail (minimum / maximum / mean)	1 / 589 / 37.8
Polygons per frame per contrail (minimum / maximum / mean)	1 / 4.5 / 1.2

underwent a two-step quality control process: a technical review made by the labelling team, followed by an expert review
 315 made by EUROCONTROL to ensure final quality. In total 4,536 h of labelling and 431 h of reviewing were performed.

4.2 Dataset Description

The GVCCS dataset Jarry et al. (2025) is the first open-access, instance-level annotated video dataset designed for contrail
 detection, segmentation, and tracking from visible ground-based sky camera imagery. It consists of 122 high-resolution video
 sequences (totalling 24,228 images) captured at the EUROCONTROL Innovation Hub in Brétigny-sur-Orge, France, using
 320 Réuniwatt’s CamVision sensor. Each sequence has been carefully annotated with temporally consistent polygon masks for
 visible contrails, including multi-instance tracking and, where possible, attribution to specific flights using aircraft trajectory
 data.

In total, the annotation team labelled 4,651 individual contrails with a total of 176,194 polygons. The sequences cover a wide
 range of durations (from 0.5 to 142.5 min per contrail), with each contrail comprising between 1 and 589 polygons (mean: 37.8).
 325 On average, each video sequence spans 96.6 min and contains approximately 193 annotated images. About 3,346 contrails are
 associated with unique flight identifiers derived from synchronized flight trajectory data filtered above 15,000 ft.

The GVCCS dataset is structured into *train/* and *test/* folders, each containing a folder with *images/*, *annotations.json* in
 COCO format, and associated flight data in parquet format. The dataset supports a range of research tasks including semantic
 and panoptic segmentation, temporal tracking, lifecycle analysis, and contrail–flight attribution, and is released under the CC
 330 BY 4.0 license.



5 Segmentation models

This section reviews the segmentation models evaluated for identifying, and for some also tracking, contrails. We focus on two model families: Mask2Former, a state-of-the-art transformer-based segmentation model, and a U-Net using a discriminative embedding loss. Both are evaluated on individual images, while only Mask2Former is additionally evaluated on videos.

335 We also explore two problem formulations: in the single-polygon case, each visible contrail fragment is treated as an independent instance; in the multi-polygon case, all fragments of a given contrail are labelled as a single instance, even if they are spatially disconnected. The single-polygon setting assumes that a subsequent linking algorithm, not implemented in this work, could later group fragments into full contrails. The multi-polygon formulation, in contrast, expects the model to infer such groupings implicitly.

340 5.1 Mask2Former

Mask2Former is a universal segmentation architecture that unifies semantic, instance, and panoptic segmentation within a single model. It is built around a hierarchical encoder-decoder structure comprising three main components: a convolutional backbone for multi-scale feature extraction, a pixel decoder that generates dense spatial embeddings, and a transformer decoder with learnable mask queries that iteratively refines segmentation predictions.

345 A central innovation in Mask2Former is its use of the so-called masked attention in the transformer decoder. Unlike standard cross-attention, which considers the entire image, masked attention limits attention to regions surrounding the current predicted masks. This localized focus enables more precise refinement of object boundaries, which is particularly beneficial for thin, high-aspect-ratio structures like contrails. The model's learnable queries act as object proposals and are refined through multiple decoding layers to generate final instance masks and class labels in an end-to-end manner.

350 An important aspect of Mask2Former's effectiveness lies in its loss function (i.e., the training objective), which guides the model to learn accurate segmentation masks and their corresponding classes. The loss function used by Mask2Former combines several components. First, it uses a classification loss that helps the model assign the correct class to each predicted mask (e.g., contrail vs. sky). Second, it includes a mask loss, which measures how closely the predicted mask matches the ground-truth mask for that object, commonly using a pixel-wise binary cross-entropy or Dice loss. Finally, Mask2Former incorporates a
355 matching step based on the Hungarian algorithm to align predictions with ground truth in an optimal, one-to-one way. This ensures that each predicted mask is evaluated against the most appropriate reference object, avoiding duplicate assignments.

A detailed technical description of the model is beyond the scope of this paper, as our focus is on applying Mask2Former to contrail segmentation; we refer the reader to the original work by Cheng et al. (2022) for a comprehensive overview of the architecture and performance on popular datasets.

360 To capture temporal dynamics inherent in contrail evolution, we extend Mask2Former to process short video sequences. Although designed for single images, the model can handle multiple consecutive frames as a 3D spatio-temporal volume by treating time as an additional axis alongside spatial dimensions, following the extension introduced by Cheng et al. (2021a).



Compared to traditional segmentation models, Mask2Former offers substantial architectural advantages. Mask R-CNN (He et al., 2017), while effective, performs detection and segmentation as separate stages, which can introduce spatial misalign-
365 ment and inefficiencies, especially when segmenting long, disconnected objects. DETR (DEtection TRansformer) (Carion et al., 2020), though end-to-end and transformer-based, primarily focuses on object detection and lacks the fine-grained spatial modelling needed for precise mask prediction. MaskFormer (Cheng et al., 2021b) introduces transformer-based decoding for segmentation, but relies on global attention, which can dilute spatial precision. Mask2Former refines this approach with masked attention and iterative refinement, leading to improved accuracy, especially in challenging tasks where objects are often thin,
370 faint, and visually ambiguous.

5.2 U-Net with Discriminative Loss

As a baseline, we implement a two steps instance segmentation model. First, we use a classical U-net architecture (Jarry et al.) for segmentation. U-Net is designed specifically for image segmentation tasks and features a symmetrical encoder-decoder structure. The encoder part of the network gradually reduces the spatial size of the input image, extracting high-level features
375 that capture the overall context. The decoder then progressively restores the spatial resolution by upsampling these features to produce a segmentation map that matches the original image size. Importantly, U-Net uses skip connections that directly link corresponding layers in the encoder and decoder. These connections allow fine-grained spatial details lost during downsampling to be recovered, improving the quality and precision of segmentation outputs.

Second, we use a similar architecture that learns a unique feature representation, or embedding, for each pixel in an image
380 by using a discriminative loss function. In this model, the final head of the U-Net does not produce a typical segmentation map with class labels. Instead, it produces an embedding for each pixel; a vector in a high-dimensional feature space. The goal is for pixels that belong to the same object instance to have similar embeddings (meaning they are close together in this feature space), while pixels belonging to different instances have embeddings that are far apart. This way, the model effectively learns to group pixels based on their learned features.

385 The process of identifying individual instances is performed in two separate steps. The first step is to generate these pixel embeddings with the U-Net, and the second step is to group or cluster these embeddings into individual instances. For clustering, we use the HDBSCAN algorithm, to find the clusters and a final k-means to associate outliers with closest cluster.

The discriminative loss used to train the model is composed of three parts. The first part, known as the pull term, encourages embeddings of pixels that belong to the same instance to be close together, making the cluster compact. The second part,
390 called the push term, forces embeddings of different instances to be sufficiently separated from each other, preventing clusters from overlapping. The third part is a regularization term that prevents the embeddings from growing too large in magnitude, which stabilizes the training process and embedding space. This combination allows the model to learn meaningful and well-separated pixel embeddings without relying on explicit object bounding boxes or pre-defined region proposals. For readers interested in the mathematical formulation and detailed rationale behind the discriminative loss, we refer to the original paper
395 by Brabandere et al. (2017).



It is important to note that this model operates only on single images. Unlike models such as Mask2Former for videos mentioned in the previous section, it does not incorporate any temporal or sequential information, nor does it include recurrent layers or mechanisms to handle videos. Extending this approach to process video sequences and incorporate temporal consistency would require significant changes to both the architecture and the algorithms used, which is outside the scope of this work.

The embedding-based approach is well suited to segmenting objects that may not be spatially continuous, such as contrails with fragmented shapes. Since the model does not require spatial continuity, it can learn to embed separate, disconnected parts of the same contrail into a similar region of the feature space if they share common visual characteristics and belong to the same label. However, this approach has its challenges. If parts of the same contrail differ significantly in appearance, due to factors like changes in lighting, atmospheric conditions, or variations in the background texture. They may be embedded differently and incorrectly assigned to separate clusters. Conversely, visually similar but unrelated contrail fragments could be mistakenly grouped together, as the model relies solely on the learned embeddings for clustering.

Figure 4 illustrates a qualitative result of the instance discriminative segmentation model. On the left, the ground truth labels are displayed, highlighting the pixel-wise assignment to contrail instances. On the right, we show the corresponding discriminative embedding space, reduced to two dimensions using Principal Component Analysis (PCA) for visualization purposes. Each point represents a pixel embedding, and colours indicate the instance it belongs to. This visualization provides insight into how the model, trained with a discriminative loss, learns to embed pixels from the same instance close together in the feature space, while separating those from different instances. The separation observed in the embedding space confirms the model's ability to cluster fragmented contrail structures, although visually similar but unrelated segments may still partially overlap in the embedding due to shared appearance features.

6 Results

This section presents the performance of the models introduced in Section 5 on contrail segmentation tasks. Our primary goal is not to achieve state-of-the-art results but to establish clear examples of application and meaningful baseline performances. By doing so, we highlight the unique opportunities offered by this dataset and provide a foundation for the research community to build upon, encouraging rapid progress in the critical field of aviation's climate impact.

6.1 Training

All models were initialised from existing pre-trained checkpoints. We trained two versions of the Mask2Former architecture for the single image segmentation task. Both models share the same core architecture but differ in the size of their transformer backbone: one uses the Swin-Base (Swin-B) configuration and the other uses the larger Swin-Large (Swin-L). The main difference between these two lies in model capacity, Swin-L has significantly more parameters, which enables it to learn richer representations at the cost of higher computational requirements.




Figure 4. The true label is displayed on the left and the discriminative embedding on the right. The latter was created using Principal Component Analysis (PCA). The colours reflect contrail instances.

Both image models were initialized from publicly available pretrained checkpoints in the Mask2Former Model Zoo¹. Each model was first pretrained on the ImageNet-21k (IN21k) (Ridnik et al., 2021) classification dataset and then fine-tuned on the COCO panoptic segmentation dataset. While COCO (Lin et al., 2014) does not include contrails, it spans a wide range of natural (including clouds and sky) and man-made objects, offering useful general-purpose segmentation features. This two-stage pre-training, IN21k followed by COCO, has been widely validated in the literature and provides a strong initialization for fine-tuning on contrail imagery.

Both the Swin-B and Swin-L variants were trained on individual image frames using 200 learnable object queries. Given our hardware setup, two NVIDIA RTX 6000 GPUs, each with 48 GB of memory, we were able to train both variants on the image dataset without significant memory limitations.

For video segmentation, we used the video-specific variant of Mask2Former, which extends the original architecture to handle temporal sequences. Like the image-based model, it also uses 200 object queries and Swin Transformer backbones, and it is initialized from a checkpoint pretrained on the YouTubeVIS 2019 dataset (Yang et al., 2019). Although YouTubeVIS does not contain contrails, its emphasis on learning temporally consistent object masks across frames makes it well suited to capture the dynamics of contrails in video data. Due to GPU memory constraints we limited both training and inference to short video clips composed of a small number of consecutive frames. While this restriction was necessary to fit within available hardware resources, particularly for memory-intensive architectures, it also shaped our training strategy. During training, these clips are randomly sampled from longer video sequences to introduce temporal diversity into the training process. By varying the starting points of the sampled clips, the model is exposed to contrails at different stages of their lifecycle, formation, elongation,

¹https://github.com/facebookresearch/Mask2Former/blob/main/MODEL_ZOO.md



445 dissipation, and in diverse atmospheric contexts. This stochastic sampling encourages the model to learn more generalizable temporal representations.

To support this setup, we trained the video Mask2Former model using both Swin-Base (Swin-B) and Swin-Large (Swin-L) backbones. However, the number of frames per clip had to be adjusted based on model capacity and memory availability. With the more lightweight Swin-B variant, we were able to train on 5-frame clips, while the higher-capacity Swin-L model
450 could only be trained on 3-frame clips due to its significantly larger memory footprint. This reflects a trade-off between temporal context and model expressiveness: longer clips may better capture the dynamic evolution of contrails, whereas larger models like Swin-L provide richer per-frame representations. Training both configurations allows us to explore how these two dimensions, temporal depth and model capacity, interact in the context of contrail segmentation.

For the U-Net model, we used a backbone based on MaxViT-B, a hybrid vision transformer architecture that combines
455 convolutional layers with self-attention mechanisms for efficient and scalable visual representation learning. This backbone was pre-trained on ImageNet-21k and subsequently fine-tuned on ImageNet-1k, providing robust feature representations to support the discriminative loss function employed during contrail segmentation training.

The training procedure for each model involved several epochs of supervised learning, with early stopping applied based on performance on a validation set.

460 We did not perform exhaustive hyper-parameter tuning for any of the models. Instead, our goal with this experimental setup was to establish baseline results and to analyse model performance both qualitatively and quantitatively under realistic computational and data constraints. All models were trained using the default hyper-parameters reported in their original publications. Tables 2 and 3 summarize the most important training parameters for each model. Note that the models differ in the specific hyper-parameters relevant to their architecture and training setup. Future work will focus on exploring more
465 sophisticated modelling strategies, systematic hyper-parameter optimization, and additional training refinements.

Remember that each model was trained and evaluated on two distinct formulations of the instance segmentation task. The first formulation treats a contrail as a single object, even if it is composed of multiple disconnected regions or fragmented segments. In this setup, the model must learn to group visually and spatially separated regions that correspond to the same physical contrail. The second task simplifies the problem by treating each visible polygon as an independent instance. In this
470 formulation, the model is not required to group disjoint segments belonging to the same contrail; instead, it simply detects and segments each distinct region. This approach corresponds to a modular processing pipeline where instance merging and flight attribution occur at a later stage, as will be discussed in future work.

6.2 Evaluation

We evaluate both semantic and instance-level segmentation performance using a combination of standard and task-adapted
475 metrics. For semantic segmentation, we report pixel-wise scores such as mean intersection over union and the Dice coefficient. For instance segmentation, we adopt the COCO evaluation protocol with modifications to better reflect the thin, elongated structure of contrails. All metrics are computed globally over the full test set. In the sections that follow, we describe our



Table 2. Default hyper-parameters for Mask2Former models.

Hyper-parameter	Default value	Notes / Differences
Training iterations	20,000	Same for image and video
Learning rate	—	$3.75e-5$ (Image), $1.25e-5$ (Video)
Batch size	—	6 (Image), 2 (Video)
Image size	$1,024 \times 1,024$	Same for image and video
Class weight	2	Same for image and video
Mask weight	5	Same for image and video
Dice weight	5	Same for image and video
Importance sample ratio	0.75	Same for image and video
Oversample ratio	3	Same for image and video
Augmentations	Rotation (90°), vertical flip, horizontal flip	Applied at image level (Image); applied at clip level (Video)

Table 3. Default hyper-parameters for the U-Net model trained with discriminative loss.

Hyper-parameter	Default value
Architecture	U-Net
Backbone	<i>tu-maxvit_base_tf_512.in1k</i>
Input image size	$1,024 \times 1,024$
Precision	16-mixed
Epochs	100
Batch size	1
Gradient accumulation steps	32
Learning rate	$5e-6$
Optimizer	AdamW (weight decay = $1e-4$)
Scheduler	Cosine with warm-up
Augmentations	Rotation (90°), vertical flip, horizontal flip

evaluation procedure, sliding window inference strategy for video models, and the rationale behind our choice of metrics. The presentation and interpretation of the results are provided at the end.

480 Model performance is evaluated using both semantic and instance-level segmentation metrics. All metrics are computed globally by aggregating predictions and ground truths across the entire test set before applying the metric calculations. This global computation prevents biases that can arise from averaging metrics computed independently on each observation (i.e., frame), which is particularly important in settings with imbalanced or sparse data such as contrail segmentation.



For semantic segmentation, we report the mean Intersection over Union (mIoU) and the Dice coefficient. The mIoU measures
485 the overlap between the predicted and ground truth binary masks by calculating the ratio of the intersection area to the union
area of the masks, thus penalizing both false positives and false negatives. The Dice coefficient, defined as twice the area of
overlap divided by the total size of the predicted and ground truth masks, emphasizes the correct overlap and is especially
sensitive to thin or fragmented structures, making it a suitable metric for evaluating contrails.

Instance segmentation performance is assessed using COCO-style metrics computed globally over the dataset. To accom-
490 modate the specific challenges posed by contrails, we adapt the IoU threshold range and denote metrics with the following
notation: $AP@[IoU\ range | size\ category | max\ detections]$, where IoU range specifies the range of IoU thresholds over which
Average Precision (AP) or Average Recall (AR) is computed, size category indicates the object size subset considered, and max
detections is the maximum number of detections per image considered. For example, $AP@[0.25:0.75 | all | 100]$ denotes the
mean average precision calculated over IoU thresholds from 0.25 to 0.75, considering all object sizes and up to 100 detections
495 per image. Object size categories follow the standard definitions used in COCO-style metrics: small objects have an area less
than 32^2 ; medium objects range between 32^2 and 96^2 pixels; large objects exceed 96^2 pixels. Metrics such as $AP@[0.25:0.75 |$
small | 100] then reflect the performance specifically on small-sized objects, under the specified IoU and detection constraints.

We restrict the IoU threshold range to [0.25, 0.75], rather than the standard COCO range of [0.50, 0.95], to better accommo-
date the elongated and thin geometry of contrails, where very high IoU thresholds are overly strict. Contrails are thin, irregular,
500 and may extend across large image portions, making exact mask overlap challenging. Under typical COCO metrics, a predic-
tion with partial but semantically correct overlap might be unfairly penalized. For example, a predicted mask overlapping only
30% of a contrail would be ignored under COCO's default minimum IoU of 0.5, but counted as a true positive under our more
lenient thresholds.

By adjusting the IoU range, the metrics better reflect practical segmentation quality for contrails, balancing sensitivity to
505 spatial accuracy with tolerance for slight misalignments and fragmentations inherent to this domain. It is important to note that
these adapted metrics are not directly comparable to standard COCO scores but are specifically tailored to provide meaningful
evaluation in the context of contrail segmentation.

This evaluation framework, combining semantic and instance segmentation metrics computed globally with appropriate IoU
thresholds and size categories, offers a comprehensive and interpretable means of assessing model performance. It facilitates
510 fair comparisons across models and supports future benchmarking on our contrail dataset.

Tables 4 and 5 summarize the results for the semantic and instance segmentation tasks, respectively. All results are reported
for both single-image and video-based models. Instance segmentation results are further disaggregated by annotation style
(multi and single-polygon). For Mask2Former models, values without parentheses correspond to the Swin-B backbone, while
those in parentheses refer to Swin-L.

515 In the semantic segmentation task, performance remains consistent across all models and variants, with Dice and mIoU
scores showing little variation. This stability is expected, as semantic segmentation only requires classifying each pixel as
either contrail or sky, without distinguishing between separate contrail instances. The U-Net model achieves results on par



Table 4. Semantic segmentation metrics. For the Mask2Former variants, values without parentheses refer to Swin-B; values in parentheses refer to Swin-L.

Metric	Single Images		Videos
	Mask2Former	U-Net	Mask2Former
Dice	0.56 (0.60)	0.59	0.57 (0.59)
mIoU	0.38 (0.43)	0.42	0.40 (0.42)

with the more advanced Mask2Former models, indicating that per-pixel contrail detection is largely driven by local visual features, such as shape, brightness, and texture, which U-Net captures effectively.

520 These results also reflect the quality and consistency of our dataset: although based on ground-level imagery, the segmentation performance is in line with results reported in previous studies using satellite data (Jarry et al.; Ortiz et al., 2025). Although differences in imaging modality and scene geometry preclude direct comparisons, the consistency in results suggests that semantic contrail segmentation is a well-posed task for modern architectures, with strong performance achievable across diverse data sources.

525 Instance segmentation results reveal clear differences between model architectures. These differences are more substantial than those observed in the semantic segmentation task, highlighting the added complexity introduced by instance-level reasoning. Mask2Former, which is designed for panoptic segmentation through object-level queries and global spatial reasoning, consistently outperforms U-Net across all instance metrics. The performance gap is particularly pronounced in the multi-polygon setting, where contrails appear fragmented and must be correctly grouped into coherent instances. These results
530 highlight the value of architectures specifically built for instance-aware tasks: Mask2Former’s ability to reason globally and associate disjoint segments makes it better suited for detecting and tracking individual contrails.

A more nuanced comparison emerges when evaluating image-based versus video-based Mask2Former models. For the Swin-B backbone, the image-based model achieves higher instance segmentation performance, while the video-based model slightly outperforms it on semantic segmentation metrics. This suggests that although video models benefit from temporal
535 consistency and motion cues, the added complexity of enforcing cross-frame coherence may introduce challenges that slightly hinder instance-level prediction accuracy, particularly when using a lower-capacity backbone like Swin-B.

In the Swin-L setting, the image-based model performs best overall. It achieves both the highest instance segmentation score and slightly superior semantic segmentation performance. These results indicate that temporal modelling does not always yield performance improvements, especially when the temporal context is limited (e.g., 3-frame clips) or when the spatial
540 representation capacity of the model is already high. The image-based model benefits from pretraining on COCO, which may favour precise spatial delineation, while the video-based variant relies on pretraining on YouTubeVIS, which is more focused on temporal coherence. However, it is important to note that the video-based model performs an additional task: tracking. By maintaining consistent instance identities across frames, it enables temporally coherent segmentation that is not achievable with



Table 5. Instance segmentation metrics. For the Mask2Former variants, values without parentheses refer to Swin-B; values in parentheses refer to Swin-L.

Type	Metric	Single Images		Videos
		Mask2Former	U-Net	Mask2Former
Multi-polygon	AP@[0.25:0.75 all 100]	0.34 (0.34)	0.05	0.31 (0.33)
	AP@[0.25:0.75 small 100]	0.21 (0.21)	0.01	0.14 (0.17)
	AP@[0.25:0.75 medium 100]	0.39 (0.40)	0.13	0.37 (0.38)
	AP@[0.25:0.75 large 100]	0.44 (0.47)	0.12	0.46 (0.47)
	AR@[0.25:0.75 all 1]	0.10 (0.10)	0.03	0.09 (0.09)
	AR@[0.25:0.75 all 10]	0.41 (0.41)	0.18	0.38 (0.40)
	AR@[0.25:0.75 all 100]	0.44 (0.44)	0.22	0.43 (0.44)
	AR@[0.25:0.75 small 100]	0.30 (0.30)	0.14	0.26 (0.29)
	AR@[0.25:0.75 medium 100]	0.50 (0.50)	0.25	0.49 (0.50)
	AR@[0.25:0.75 large 100]	0.55 (0.55)	0.22	0.57 (0.56)
Single-polygon	AP@[0.25:0.75 all 100]	0.35 (0.37)	0.06	0.31 (0.34)
	AP@[0.25:0.75 small 100]	0.24 (0.26)	0.03	0.17 (0.21)
	AP@[0.25:0.75 medium 100]	0.44 (0.45)	0.14	0.41 (0.43)
	AP@[0.25:0.75 large 100]	0.37 (0.43)	0.11	0.46 (0.47)
	AR@[0.25:0.75 all 1]	0.08 (0.08)	0.03	0.07 (0.08)
	AR@[0.25:0.75 all 10]	0.37 (0.38)	0.18	0.35 (0.37)
	AR@[0.25:0.75 all 100]	0.44 (0.45)	0.21	0.42 (0.45)
	AR@[0.25:0.75 small 100]	0.33 (0.34)	0.15	0.28 (0.32)
	AR@[0.25:0.75 medium 100]	0.53 (0.53)	0.26	0.52 (0.55)
	AR@[0.25:0.75 large 100]	0.54 (0.56)	0.25	0.58 (0.60)

image-based models. All in all, the metrics reported here are computed on a per-frame basis and do not account for flickering or instance identity consistency over time. These temporal aspects are particularly important in video applications and are not captured by the conventional frame-level evaluation scores presented herein.

Overall, Swin-L outperforms Swin-B across all setups, reinforcing the benefit of increased model capacity for fine-grained spatial understanding and instance-level reasoning. Nonetheless, this comes at the cost of higher computational requirements, particularly in the video setting, underscoring a trade-off between performance and scalability.

Another important trend observed in the evaluation is that model performance is strongly influenced by contrail size and detection caps. Generally speaking, larger contrails are segmented more accurately due to their higher pixel counts and lower ambiguity, while allowing more predicted instances (e.g., increasing the detection limit) improves recall by removing constraints



on how many objects can be reported. These trends are consistent with general findings in object detection and reinforce the shared challenges between contrail segmentation and broader instance segmentation tasks.

555 Comparing the multi-polygon and single-polygon formulations reveals a difference in task difficulty: the single-polygon setting is inherently easier. Across all models and data modalities, instance segmentation metrics are consistently higher when using the single-polygon formulation. This is because the task removes the need to group fragmented or spatially disjoint contrail segments into separate instances. Instead, all parts of a contrail, regardless of their separation, are treated as a single mask, greatly simplifying the model's objective. The model is no longer required to learn complex grouping strategies or
560 reason over spatial and temporal discontinuities. Note that semantic segmentation metrics remain virtually unchanged between the two formulations, indicating that identifying contrail pixels is equally feasible in both cases. The difference lies solely in how those pixels are grouped into instances. This distinction confirms that the main challenge in the multi-polygon task is not pixel classification but instance association.

These results have important practical implications for different contrail detection scenarios. For older contrails, such as
565 those typically observed in satellite imagery or in ground-based images when the contrail formed outside the camera's field of view, it is extremely difficult to associate the contrail with its source flight. In these cases, the only viable option is to group visible fragments into instances based solely on visual information. This makes multi-polygon instance segmentation essential, as it allows models to detect and associate disjoint contrail segments without relying on external data. Our dataset and Mask2Former-based models are specifically designed for this setting, enabling effective instance-level detection even when
570 contrails are fragmented, occluded, or spatially disconnected.

In contrast, when a contrail forms directly above the camera and additional data such as aircraft trajectories and wind fields are available, a different approach becomes feasible. In these situations, one can perform single-polygon instance segmentation, where contrail fragments are grouped into a single instance using post-hoc association based on flight paths and advection. This formulation is simpler from a computer vision perspective and is commonly used in the literature (Ortiz et al., 2025; Chevallier
575 et al., 2023; Van Huffel et al., 2025), mainly because multi-polygon annotated datasets have not been available until now. However, this method depends on access to external data and is only applicable to contrails formed during the observation window, after the aircraft has entered the scene.

By supporting both the multi- and single-polygon formulations, our dataset enables training and evaluation across a broader set of operational use cases. The multi-polygon task is essential for vision-only detection of older contrails or those in satel-
580 lite imagery, while the single-polygon formulation may be more suitable when additional metadata enables contrail-to-flight attribution. This distinction will be further explored in future work focused on linking contrails to their source aircraft.

6.3 Illustrative examples

We present two test-set examples to illustrate the challenges of the multi-polygon contrail segmentation task. In both cases, we compare predictions from image-based and video-based versions of the Mask2Former model, trained from pretrained Swin-
585 L backbones. These examples highlight how temporal context affects instance predictions and expose typical failure modes, including contrail fragmentation, occlusion by clouds, and confusion between contrails and visually similar cloud structures.



Figure 5 shows a frame from 25 April 2024 at 05:51:00, under clear-sky conditions. The background is uniformly blue, providing favourable conditions for both human and machine segmentation. The corresponding ground-truth annotations include several contrails labelled as fragmented (e.g., identifiers 0, 1, and 5), based on known flight trajectories available to annotators during the labelling process. This makes the example suitable for evaluating instance-level understanding in the multi-polygon setting.



(a) Raw image.

(b) Ground truth annotations.

Figure 5. Raw image and ground truth annotations for 25 April 2024 at 05:51:00.

Despite the favourable background, both models exhibit instance-level errors. The image-based model correctly infers that contrail 1 is fragmented, but detects just one segment of contrail 0, missing the other entirely. It completely misses contrail 4 and erroneously merges contrails 5 and 6 into a single prediction. The video-based model makes similar mistakes: it also merges contrails 5 and 6, and fails to detect contrail 4. Additionally, it predicts the second fragment of contrail 0 but assigns it to a different instance, and it incorrectly splits contrail 1 into two separate instances.

From a semantic segmentation perspective, both models perform relatively well, as expected in a high-contrast scene. The image-based model achieves a Dice score of 0.76 and a mean IoU of 0.64, while the video-based model slightly outperforms it with a Dice of 0.79 and mean IoU of 0.67. However, due to the instance grouping errors, the image model achieves a slightly higher AP@[0.25:0.75 | all | 100] (0.62) than the video model (0.55).

Figure 7 shows a more challenging frame captured on 19 November 2023 at 08:49:30. Here, several cirrus clouds are present in the background, which introduces ambiguity, as some of these cloud structures resemble contrails. This scene also includes multiple contrails that are spatially aligned and fragmented, increasing the complexity of the instance segmentation task.

This scene illustrates a common failure mode: fragmentation and miss-grouping of visually aligned but semantically distinct contrails. Contrail 6 is split into two segments with contrail 0 lying in between; although they appear collinear, contrail 0 is




Figure 6. Predicted instances for the frame shown in Fig. 5, using Swin-L models with image and video inputs.




Figure 7. Raw image and ground truth annotations for 19 November 2023 at 08:49:30.

a distinct instance generated by a separate flight. Contrail 7 appears shortly after and may be miss-associated with contrails 6 and 0 in the absence of flight metadata. The image-based model correctly separates contrail 0 from 6, but incorrectly merges




Figure 8. Predicted instances for the frame shown in Fig. 7, using Swin-L models with image and video inputs.

contrails 6 and 7. The video model groups all three, 6, 0, and 7, into a single prediction. Interestingly, this error reflects a plausible human interpretation without flight context, highlighting the challenge of the task.

Both models fail to detect contrails 1 and 8, which are partially occluded by clouds. They also produce a false positive (labelled as contrail 9), segmenting a cirrus structure that resembles a contrail. While the dataset is of high quality and was carefully annotated with access to flight information, some visually ambiguous cases, such as the one discussed, remain inherently difficult to label with certainty. In this example, the predicted region resembles a contrail in both structure and intensity, making it unclear whether the false positive stems from a model error or an understandable omission in the ground truth. These rare edge cases highlight the potential influence of mild label noise in visually complex scenes. Future work could benefit from complementary strategies such as confident learning (Northcutt et al., 2021) to further refine annotations and improve robustness in borderline cases.

Semantic segmentation performance in this scene is lower than in the previous one, reflecting increased difficulty. The image model achieves a Dice score of 0.61 and mIoU of 0.43, while the video model scores 0.70 and 0.54, respectively. Instance-level AP@[0.25:0.75 | all | 100]s are 0.35 and 0.37, respectively, similar to the average metrics, making this a representative case.

These examples illustrate several key challenges in multi-polygon contrail segmentation: (1) correct grouping of fragmented contrail segments from the same flight; (2) visual ambiguity due to clouds that resemble contrails; (3) occlusion; and (4) spatial overlap of contrails from different flights. While video-based models benefit from temporal information, they may over-group distinct instances. Image-based models avoid this but often fail to connect fragmented segments. Overall, these examples demonstrate the inherent difficulty of the task and the limitations of current models.



7 Conclusions

This work introduces a new dataset (Jarry et al., 2025) and baseline models for contrail segmentation from ground-based camera imagery. Our experiments show that modern computer vision methods, particularly panoptic segmentation models like Mask2Former, can be effectively applied to this task, especially when using large pre-trained models and temporal information.

630 However, performance gains often come at the cost of increased computational and memory demands, highlighting a trade-off between accuracy and practicality.

The main contribution of this study is the release of the first video annotated dataset specifically designed for instance-level contrail segmentation, tracking and flight attribution in the visual spectrum. Along with detailed evaluation metrics, including average precision and recall across multiple intersection-over-union thresholds and object size bins, this benchmark provides a
635 reproducible baseline for further research in this emerging field.

A key limitation of our current setup is that the visible-light camera restricts observations to daytime conditions. Yet contrails often have their greatest radiative impact at night, when they trap outgoing long wave radiation and contribute to atmospheric warming. To address this, we are deployed a co-located infrared imaging system that enables continuous, day-and-night monitoring. This may also allow us to begin estimating the radiative forcing of individual contrails under real atmospheric condi-
640 tions.

In parallel, we are working on a contrail-to-flight attribution algorithm that links observed contrails to specific aircraft using automatic dependent surveillance–broadcast (ADS-B) trajectory data. This tool, and the associated data and code, will be openly released in a future publication. Attribution is of utmost importance because it allows each contrail to be linked to detailed aircraft and engine parameters, such as aircraft type, engine model, fuel burn rate, flight altitude, and ambient
645 conditions. These inputs are necessary to reproduce the contrail using physical models like CoCiP, assess its expected properties (e.g., ice crystal number, optical depth, lifetime), and ultimately validate or refine these models using real-world observations.

We are also extending this work by annotating a new dataset of contrails in satellite imagery, with instance-level and sequence-based labels. This dataset will allow us to test and evaluate the full multi-scale tracking pipeline proposed in this paper: starting from high-resolution, ground-based detection, followed by attribution to flights, and finally linking to the same
650 contrails as they evolve in satellite imagery. This approach offers a unique opportunity to study contrail formation, spreading, and dissipation over time and at scale. We also plan to use our ground-based dataset to evaluate the predictions of physical models such as CoCiP. Direct comparisons between observed and simulated contrail evolution will help assess model accuracy and potentially inform improvements in contrail forecasting and climate modelling.

Ideally, contrail detection, tracking, and attribution should be addressed by a single deep learning architecture capable of
655 jointly processing video, flight trajectory data, and meteorological fields. A model such as a variant of Mask2Former could be adapted for this purpose. Integrating these tasks into one architecture would enable end-to-end learning and exploit the complementary nature of the inputs, weather conditions and aircraft traffic data are highly informative for both detecting and tracking contrails. However, this integration is not straightforward. It requires careful design of input data representations to handle spatio-temporal and multi-modal inputs, the creation of aligned and consistent annotations for all tasks, and the



660 development of loss functions that balance competing objectives across detection, segmentation, tracking, and attribution. Despite these challenges, we encourage the research community to explore this unified approach.

More broadly, we hope this work encourages the development of similar ground-based contrail monitoring systems in other regions. A collaborative, open-science approach, sharing datasets, models, and observational infrastructure, will be essential to building a geographically diverse and temporally continuous picture of contrail behaviour. We view this paper as a first step
665 toward a data-driven ecosystem for contrail research: one that integrates physical modelling with observational data, spans spatial and temporal scales, and supports long-term efforts to better understand and reduce aviation's impact on the climate.

8 Data availability

The dataset used and analysed during this study is available in the GVCCS: Ground Visible Camera Contrail Sequences repository and can be accessed at <https://doi.org/10.5281/zenodo.16419651>. The recommended citation for this dataset is:
670 Jarry, G., Very, P., Ballerini, F., and Dalmau, R.: GVCCS : Ground Visible Camera Contrail Sequences, <https://doi.org/10.5281/zenodo.16419651>, 2025. The dataset is openly available under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. No access restrictions apply. All data necessary to replicate the results in this article are included.

Author contributions. G. Jarry: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Software, Visualization, Supervision, Writing (original draft preparation); R. Dalmau: Formal analysis, Investigation, Methodology,
675 Software, Visualization, Writing (original draft preparation); P. Very: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Software, Visualization, Supervision, Writing (original draft preparation); F. Ballerini: Data curation, Project administration, Writing (review and editing); S-D. Bocu: Data curation, Formal analysis, Software, Writing (review and editing).

Competing interests. The authors declare that they have no conflict of interest.

680 *Acknowledgements.* We would like to sincerely thank Dr. Alexe Bogdan for the insightful discussions and valuable proposals on the discriminative loss. Our heartfelt appreciation also goes to our manager, Robin Deransy, for the continuous support and encouragement throughout the project. We are grateful to Réuniwatt for their support related to the ground-based sky imaging systems. Additionally, we gratefully acknowledge Encord for providing the labelling platform and workforce that made our data annotation possible.



References

- 685 Appleman, H.: The Formation of Exhaust Condensation Trails by Jet Aircraft, *Bulletin of the American Meteorological Society*, 34, 14 – 20,
<https://doi.org/10.1175/1520-0477-34.1.14>, 1953.
- Borella, A., Boucher, O., Shine, K. P., Stettler, M., Tanaka, K., Teoh, R., and Bellouin, N.: The importance of an informed choice of CO
 2-equivalence metrics for contrail avoidance, *Atmospheric Chemistry and Physics*, 24, 9401–9417, 2024.
- Brabandere, B. D., Neven, D., and Gool, L. V.: Semantic Instance Segmentation with a Discriminative Loss Function, *CoRR*, abs/1708.02551,
 690 <http://arxiv.org/abs/1708.02551>, 2017.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S.: End-to-End Object Detection with Transformers, in:
 European Conference on Computer Vision (ECCV), pp. 213–229, 2020.
- Chen, L.-C., Papandreou, G., Schroff, F., and Adam, H.: Rethinking atrous convolution for semantic image segmentation, *arXiv preprint*
arXiv:1706.05587, 2017.
- 695 Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H.: Encoder-decoder with atrous separable convolution for semantic image
 segmentation, in: *Proceedings of the European conference on computer vision (ECCV)*, pp. 801–818, 2018.
- Cheng, B., Choudhuri, A., Misra, I., Kirillov, A., Girdhar, R., and Schwing, A. G.: Mask2Former for Video Instance Segmentation, *arXiv*
preprint arXiv:2112.10764, extends Mask2Former to video by processing 3D segmentation volumes without modifying architecture or
 loss, 2021a.
- 700 Cheng, B., Schwing, A. G., and Kirillov, A.: Per-Pixel Classification is Not All You Need for Semantic Segmentation, in: *Advances in Neural*
Information Processing Systems (NeurIPS), 2021b.
- Cheng, B., Misra, I., Schwing, A. G., Kirillov, A., and Girdhar, R.: Masked-Attention Mask Transformer for Universal Image Segmentation,
 in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Chevallier, R., Shapiro, M., Engberg, Z., Soler, M., and Delahaye, D.: Linear contrails detection, tracking and matching with aircraft using
 705 geostationary satellite and air traffic data, *Aerospace*, 10, 578, 2023.
- Dai, Z., Liu, H., Le, Q. V., and Tan, M.: Coatnet: Marrying convolution and attention for all data sizes, *Advances in neural information*
processing systems, 34, 3965–3977, 2021.
- Duda, D. P., Minnis, P., Khlopenkov, K., Chee, T. L., and Boeke, R.: Estimation of 2006 Northern Hemisphere contrail coverage using
 MODIS data, *Geophysical Research Letters*, 40, 612–617, 2013.
- 710 Forkert, T., Strauss, B., and Wendling, P.: A new algorithm for the automated detection of jet contrails from NOAA-AVHRR satellite images,
 in: *Proc. of the 6th AVHRR Data Users' Meeting*, pp. 513–519, EUMETSAT-Joint Research Centre of the Commission of the EC, 1993.
- Fritz, T. M., Eastham, S. D., Speth, R. L., and Barrett, S. R.: The role of plume-scale processes in long-term impacts of aircraft emissions,
Atmospheric Chemistry and Physics, 20, 5697–5727, 2020.
- Geraedts, S., Brand, E., Dean, T. R., Eastham, S., Elkin, C., Engberg, Z., Hager, U., Langmore, I., McCloskey, K., Ng, J. Y.-H., et al.: A
 715 scalable system to measure contrail formation on a per-flight basis, *Environmental Research Communications*, 6, 015 008, 2024.
- Gierens, K., Matthes, S., and Rohs, S.: How Well Can Persistent Contrails Be Predicted?, *Aerospace*, 7,
<https://doi.org/10.3390/aerospace7120169>, 2020.
- Gourgue, N., Boucher, O., and Barthès, L.: A dataset of annotated ground-based images for the development of contrail detection algorithms,
Data in Brief, 59, 111 364, 2025.



- 720 He, K., Gkioxari, G., Dollár, P., and Girshick, R.: Mask R-CNN, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 2961–2969, 2017.
- Jarry, G., Very, P., Heffar, A., and Torjman-Levavasseur, V.: Deep Semantic Contrails Segmentation of GOES-16 Satellite Images: A Hyperparameter Exploration.
- Jarry, G., Very, P., Ballerini, F., and Dalmau, R.: GVCCS : Ground Visible Camera Contrail Sequences,
725 <https://doi.org/10.5281/zenodo.16419651>, 2025.
- Kulik, L.: Satellite-based detection of contrails using deep learning, Ph.D. thesis, Massachusetts Institute of Technology, 2019.
- Lee, D. S., Fahey, D. W., Skowron, A., Allen, M. R., Burkhardt, U., Chen, Q., Doherty, S. J., Freeman, S., Forster, P. M., Fuglestedt, J., et al.: The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018, *Atmospheric Environment*, 244, 117 834, 2021.
- 730 Lin, T., Maire, M., Belongie, S. J., Bourdev, L. D., Girshick, R. B., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L.: Microsoft COCO: Common Objects in Context, *CoRR*, abs/1405.0312, 2014.
- Low, J., Teoh, R., Ponsonby, J., Gryspeerdt, E., Shapiro, M., and Stettler, M. E.: Ground-based contrail observations: comparisons with reanalysis weather data and contrail model simulations, *Atmospheric Measurement Techniques*, 18, 37–56, 2025.
- Mannstein, H., Meyer, R., and Wendling, P.: Operational detection of contrails from NOAA-AVHRR-data, *International Journal of Remote*
735 *Sensing*, 20, 1641–1660, 1999.
- McCloskey, K. J. F., Geraedts, S. D., Jackman, B. H., Meijer, V. R., Brand, E. W., Fork, D. K., Platt, J. C., Elkin, C., and Van Arsdale, C. H.: A Human-Labeled Landsat-8 Contrails Dataset, in: Proceedings of the ICML Workshop on Tackling Climate Change with Machine Learning, 2021.
- Meijer, V. R., Kulik, L., Eastham, S. D., Allroggen, F., Speth, R. L., Karaman, S., and Barrett, S. R.: Contrail coverage over the United States
740 before and during the COVID-19 pandemic, *Environmental Research Letters*, 17, 034 039, 2022.
- Meijer, V. R., Eastham, S. D., Waitz, I. A., and Barrett, S. R.: Contrail altitude estimation using GOES-16 ABI data and deep learning, *Atmospheric Measurement Techniques*, 17, 6145–6162, 2024.
- Ng, J. Y.-H., McCloskey, K., Cui, J., Meijer, V. R., Brand, E., Sarna, A., Goyal, N., Van Arsdale, C., and Geraedts, S.: Opencontrails: Benchmarking contrail detection on goes-16 abi, *arXiv preprint arXiv:2304.02122*, 2023.
- 745 Northcutt, C., Jiang, L., and Chuang, I.: Confident Learning: Estimating Uncertainty in Dataset Labels, *J. Artif. Int. Res.*, 70, 1373–1411, <https://doi.org/10.1613/jair.1.12125>, 2021.
- Ortiz, I., García-Heras, J., Jafarimoghaddam, A., and Soler, M.: Enhancing GOES-16 Contrail Segmentation through Ensemble Neural Network Modeling and Optical Flow Corrections, *Authorea Preprints*, 2025.
- Pratt, W. K.: Digital image processing: PIKS Scientific inside, vol. 4, Wiley Online Library, 2007.
- 750 Ridnik, T., Baruch, E. B., Noy, A., and Zelnik-Manor, L.: ImageNet-21K Pretraining for the Masses, *CoRR*, abs/2104.10972, 2021.
- Riggi-Carolo, E., Dubot, T., Sarrat, C., and Bedouet, J.: AI-Driven Identification of Contrail Sources: Integrating Satellite Observation and Air Traffic Data, *Journal of Open Aviation Science*, 1, 2023.
- Ronneberger, O., Fischer, P., and Brox, T.: U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, pp. 234–241, Springer, 2015.
- 755 Sarna, A., Meijer, V., Chevallier, R., Duncan, A., McConaughay, K., Geraedts, S., and McCloskey, K.: Benchmarking and improving algorithms for attributing satellite-observed contrails to flights, *EGUsphere*, 2025, 1–58, 2025.



- Schumann, U.: A contrail cirrus prediction model, *Geoscientific Model Development*, 5, 543–580, <https://doi.org/10.5194/gmd-5-543-2012>, 2012.
- 760 Schumann, U., Hempel, R., Flentje, H., Garhammer, M., Graf, K., Kox, S., Lösslein, H., and Mayer, B.: Contrail study with ground-based cameras, *Atmospheric Measurement Techniques*, 6, 3597–3612, 2013.
- Sun, J. and Roosenbrand, E.: Few-Shot Contrail Segmentation in Remote Sensing Imagery With Loss Function in Hough Space, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2025.
- Teoh, R., Schumann, U., and Stettler, M. E.: Beyond contrail avoidance: Efficacy of flight altitude changes to minimise contrail climate forcing, *Aerospace*, 7, 121, 2020.
- 770 Teoh, R., Engberg, Z., Schumann, U., Voigt, C., Shapiro, M., Rohs, S., and Stettler, M.: Global aviation contrail climate effects from 2019 to 2021, *EGUsphere*, 2023, 1–32, 2023.
- Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., and Li, Y.: Maxvit: Multi-axis vision transformer, in: *European conference on computer vision*, pp. 459–479, Springer, 2022.
- Van Huffel, J., Ehrmanntraut, R., and Croes, D.: Contrail Detection and Classification using Computer Vision with Ground-Based Cameras, in: *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, pp. 1–6, IEEE, 2025.
- Vazquez-Navarro, M., Mannstein, H., and Mayer, B.: An automatic contrail tracking algorithm, *Atmospheric Measurement Techniques*, 3, 1089–1101, 2010.
- Yang, L., Fan, Y., and Xu, N.: Video instance segmentation, in: *ICCV*, 2019.