



# A machine learning approach for contrail detection and persistence prediction using airborne measurements from the ECLIF II/ND-MAX and ecoDemonstrator flight campaigns

Ariadne K. Papamichou<sup>1</sup>, Evanthia Kallou<sup>2</sup>, Richard H. Moore<sup>3</sup>, Holger Pfaender<sup>4</sup>, and Dimitri N. Mavris<sup>5</sup>

<sup>1,2,4,5</sup>Georgia Institute of Technology, Atlanta, GA, USA

<sup>3</sup>NASA Langley Research Center, Hampton, VA, USA

**Correspondence:** Evanthia Kallou (ekallou3@gmail.com)

**Abstract.** Contrails or condensation trails are a major contributor to aviation-induced cloudiness, which represents a significant, yet highly uncertain, component of aviation's environmental impact. The reliable detection and characterization of contrails has become increasingly important for quantifying their radiative forcing and developing mitigation strategies. However, contrail detection and prediction remains challenging due to their variable optical properties, lack of accuracy in humidity, temperature and pressure sensor measurements, as well as in weather prediction. This paper investigates the use of machine vision on board aircraft to inform contrail formation and classification in real time. The focus of this work is the comparison of two major airborne measurement campaigns conducted by NASA, industry, and international partners: the 2018 NASA - DLR ECLIF II/ND-MAX and the 2023 NASA - Boeing ecoDemonstrator flight tests. The atmospheric data from the campaigns are used with the Schmidt-Appleman criterion to identify periods of contrail formation. For the ND-MAX dataset classification accuracies of 94.0 %, 87.1 %, and 81.5 % were obtained for contrail absence, short-lived and persistent categories, respectively. For the ecoDemonstrator dataset the corresponding accuracies were 90.4 %, 89.5 %, and 93.3 %. Comparison between the two campaigns reveals that camera placement affects the classification performance; longer visible contrail segments improve detection for the absence and short-lived categories while reduced airframe intrusion in the camera's field of view improves persistent contrail classification. Based on these findings, some recommendations for camera placement on flights are provided. The approach requires only an onboard camera as additional instrumentation, making it a cost-effective and scalable tool that may complement existing contrail monitoring and mitigation strategies.

## 1 Introduction

Condensation trails (contrails) are line-shaped clouds of ice particles that can persist for several hours under ice supersaturated conditions and may subsequently spread into climate-altering contrail cirrus (Kärcher, 2018). Current assessments, such as those by Lee et al. (2021), indicate that in the cumulative impact of aviation's CO<sub>2</sub> emissions from 1940 to 2018 contributed roughly +34.3 (28, 40) mW m<sup>-2</sup> of effective radiative forcing (ERF) on Earth's climate system, while much shorter-lived contrail cirrus contributed an even greater +57.4 (17, 98) mW m<sup>-2</sup>, where parentheses indicate the 10% and 90% likelihood

ranges. Despite the much larger uncertainty associated with the contrail cirrus ERF relative to the carbon dioxide ERF, it is clear that contrail cirrus exert a significant warming impact. ERF, which accounts for rapid atmospheric adjustments in temperatures and cloud cover, is the most comprehensive metric to quantify such climate impacts. Current analyses suggest aviation accounts for approximately 3.5 % of total human-induced climate forcing when both CO<sub>2</sub> and non-CO<sub>2</sub> emissions are considered (Kärcher, 2018; Lee et al., 2021). The aviation sector's role in global anthropogenic climate change has become increasingly prominent, with growing air traffic intensifying the pressure.

Notably, while the average ERF is higher for contrail cirrus than for CO<sub>2</sub> emissions, the quantification of the former is subject to a substantially higher uncertainty. This uncertainty for the ERF, which approaches 70 % compared to 17 % for CO<sub>2</sub> (Lee et al., 2021), arises from factors such as the incomplete understanding of contrail evolution and dispersion (Singh et al., 2024), its short lifespan compared to CO<sub>2</sub> (International Air Transport Association, 2024), the limited available observational studies and modeling frameworks and the inconsistencies among radiation schemes in climate models, among others (National Academies of Sciences, Engineering, and Medicine, 2025).

The short atmospheric residence of contrails makes them an attractive target for mitigation, as any operational or technological intervention that could lead to a reduction or elimination of contrails will result in an immediate reduction in ERF. Multiple flight campaigns have demonstrated how the use of sustainable aviation fuels (SAF) can help reduce contrail cloudiness. Reducing soot particle emissions leads to the formation of fewer ice crystals and contrails with lower optical depth, and therefore achieves a significant reduction in the climate impact of contrails (Voigt et al., 2021).

The scale of aviation's climate impact, as well as the urgent need to reduce it, has prompted regulatory action, as well as calls for comprehensive research agendas (National Academies of Sciences, Engineering, and Medicine, 2025). The European Union has recently mandated the systematic monitoring and reporting of all non-CO<sub>2</sub> emissions in the aviation sector (European Commission, Directorate-General for Climate Action, 2024). In turn, this regulatory requirement creates a direct imperative for the development of robust methodologies that are capable of accurately collecting, characterizing, and quantifying contrail formation and its associated climate impacts.

Contrails form when water vapor from the hot engine exhaust mixes with the cold and humid surrounding air. If the cooling plume becomes supersaturated with respect to liquid water (i.e., %RH<sub>water</sub> > 100%), liquid cloud droplets form and subsequently freeze homogeneously. If the surrounding air is supersaturated with respect to ice (i.e., %RH<sub>ice</sub> > 100%), the newly formed ice crystals persist and grow. Otherwise, the ice crystals sublimate. Multiple factors influence the contrail formation, including the local atmospheric temperature, pressure, and relative humidity, as well as the amount of heat and water vapor emitted by the engine (Schumann, 1996). Depending on those conditions, contrails exhibit behavior that can broadly be classified into two main categories: short-lived linear contrails that occur in regions where %RH<sub>ice</sub> < 100% and persistent contrails or contrail cirrus that occur in ice supersaturated regions (ISSRs) where %RH<sub>ice</sub> > 100%.

Regional analyses of contrail effects, such as the one presented by Teoh et al. (2022) over the North Atlantic aviation corridors have shown significant seasonal and diurnal variability in both contrail formation and persistence. The ability to estimate contrail impact using a flight-by-flight approach shows the value of more accurate spatial and temporal predictions in high traffic regions.

## 1.1 Contrail detection and prediction methods

Contrail detection and prediction are being pursued with several methods in operation or under active development, from  
60 satellite remote sensing to ground based observations and computational methods. One of the approaches relies on numerical  
predictions from atmospheric models, such as from the European Center for Medium-Range Weather Forecasts (ECMWF)  
and the Global Forecast System (GFS) model. By combining satellite observations with in situ measurements, these models  
estimate meteorological conditions along the flight path. The resulting atmospheric conditions along the aircraft route can be  
used to test the Schmidt–Appleman criterion (SAc) for contrail formation (Appleman, 1953). Because it relies on existing  
65 forecast data, an important advantage is that it does not require modification of aircraft instrumentation. Unfortunately, nu-  
merical weather prediction models are limited by inaccurate estimates of humidity at cruise altitudes, due to the scarcity of  
in situ humidity measurements in the upper troposphere, thus leading to significant uncertainty in the prediction of contrail  
persistence (Hofer et al., 2024).

In order to model the full lifecycle of contrails, from initial formation conditions till dispersion, physics-based simulation  
70 tools such as the Contrail Cirrus Prediction model (CoCiP; Schumann, 2012) and the Aircraft Plume Chemistry, Emissions,  
and Microphysics Model (APCEMM; Fritz et al., 2020) have been developed throughout the years. Both models capture the  
microphysical properties of contrails, at different levels of fidelity by integrating aircraft emission parameters with background  
atmospheric conditions. More specifically, CoCiP uses a Lagrangian-Gaussian plume model which simulates the contrail cross-  
section as a descending Gaussian plume. Due to the fact that mixing and bulk cloud processes are represented within CoCiP  
75 using semi-analytical formulations, it remains more computationally efficient, and its open-source availability has facilitated  
its broad adoption across the research community (Akhtar Martínez et al., 2025). On the other hand, APCEMM is more com-  
putationally expensive, since it uses a two-dimensional gridded approach to model the plume chemistry and microphysics, and  
in this way provides a more detailed representation of contrail evolution compared to the parametrized Gaussian formulation  
within CoCiP.

80 Ground-based approaches to contrail detection and prediction have been explored. In their work, Schumann et al. (2013)  
apply photogrammetric methods for a use case of two different types of wide-angle video cameras, with the intent to improve  
the collection of ground-based observations regarding contrail geometric motion speeds, altitudes and widths. In another study,  
Schumann et al. (2017) created a database with microphysical and geometric contrail properties using in situ past contrail  
measurements with optical particle spectrometers, visual observations and remote sensing data from satellites, cameras as well  
85 as ground-based and airborne lidar observations, for individual aircraft contrails. This database enables the testing of contrail  
models and thus, improved understanding of contrail formation, observed properties and lifespan evolution. More recently, Low  
et al. (2025) established a framework to analyze contrail formation, persistence and geometric properties using ground-based  
video footage and then validated those findings against CoCiP simulations with actual flight trajectories.

Lidar have been used in research settings, both from ground-based installations and airborne platforms to characterize con-  
90 trail optical and microphysical properties. Ground-based scanning lidar instruments, as demonstrated in the work by Freuden-  
thaler et al. (1995), can retrieve contrail geometry and measure its spatial growth. Airborne lidar systems use laser pulses to



probe the optical properties of the atmosphere, retrieving depolarization ratios and backscatter coefficients of ice crystals to enable the assessment of contrail optical properties. Recent work by Soleimanpour et al. (2026) leverages airborne lidar observations from the ML-CIRRUS and CIRRUS-HL campaigns, for development and verification, respectively, along with aircraft location data, to detect the occurrence of contrails within existing cirrus clouds. Mandija et al. (2026) developed an approach that combines Raman lidar, Automatic Dependent Surveillance-Broadcast (ADS-B) flight data and ECMWF ERA 5 reanalysis data to improve the nighttime detection and characterization of individual aircraft contrails. However, despite their promise in terms of high vertical resolution, onboard lidar systems present significant certification, economic, and operational challenges that render fleet-wide implementation difficult in the near term.

100 Satellite-based detection is currently one of the preferred approaches for large-scale contrail monitoring. Geraedts et al. (2024) demonstrated a scalable system that detects contrails in satellite imagery (GOES-16 Advanced Baseline Imager (ABI) imagery) using a computer vision algorithm and then matches the observed contrails to specific flights. Convolutional neural network (CNN) for contrail segmentation in satellite imagery were used in (Hoffman et al., 2023) to detect contrails, also in GOES-16 ABI imagery. Deep learning techniques for contrail altitude estimation (Meijer et al., 2024) can be combined with existing contrail detection methods to provide 3D estimates of contrail locations, using GOES-16 ABI satellite data across multiple imaging bands. Riggi-Carolo et al. (2023) developed an algorithm to enable automatic contrail detection using satellite imagery. Their methodology uses air traffic and meteorological data along with contrail imagery, to match contrails to the respective aircraft trajectories. Another example of satellite imagery use for contrail detection, currently in development, is Project Contrails by Google Research (Google Research, 2023; Sarna et al., 2025; Ng et al., 2023; Geraedts et al., 2024). Project Contrails aims to predict the timing and location of contrail formation by integrating satellite imagery with meteorological and flight trajectory data. In addition, it uses computer vision to identify contrails and associate them with individual flights. A major strength of this approach is its potential to make accurate near real-time detection and predictions on regions where contrails might form. At the same time, this method depends on several datasets, which can require high computational resources during training and deployment.

## 115 1.2 Current Scope and Contributions

The approaches summarized above offer several benefits but also involve practical constraints that have shaped the present work. The aim of this study is to investigate a method onboard aircraft for detecting contrails in real-time and estimating their persistence without excessive computational demands or complex fleet-wide implementation requirements. The analysis was based on publicly available data from the 2018 NASA-DLR ECLIF II/ND-MAX and the 2023 NASA-Boeing ecoDemonstrator flight tests, which combine airborne atmospheric measurements with corresponding aircraft-mounted visual footage (NASA Airborne Science Program; NASA Aeronautics Mission Directorate). The atmospheric datasets provided the reference conditions needed to assess whether the surrounding environment was conducive to contrail formation and persistence, and they supplied the labels used to train the model to classify contrails in video footage. Together, these datasets enable an evaluation of machine vision as a means of real-time contrail detection and persistence prediction onboard aircraft. A prototype of the model described in the following sections was tested in (Papamichou et al., 2026) with promising results, particularly in cases

where no contrails were formed. The detection of short-lived and persistent contrails in the model was significantly higher than random classification in these two identified categories, at 81.3 % and 77.5 %, respectively, but with room for improvement. This led to the investigation and comparison of the impact of camera placement on the machine learning model between the 2018 NASA - DLR ECLIF II/ND-MAX and the 2023 NASA - Boeing ecoDemonstrator flight measurement campaigns. The results of these two campaigns are compared in Sect. 3 to identify possible areas of improvement in camera placement.

The paper is organized into the following sections. The methodology presented in Section 2 will outline the machine vision algorithm used to detect contrail formation. In Sect. 3, the results of the proposed methodology will be demonstrated on each measurement campaign, and in Sect. 4, the results obtained will be discussed in detail. Finally, Section 5 will summarize the main contributions of the study and outline directions for future research.

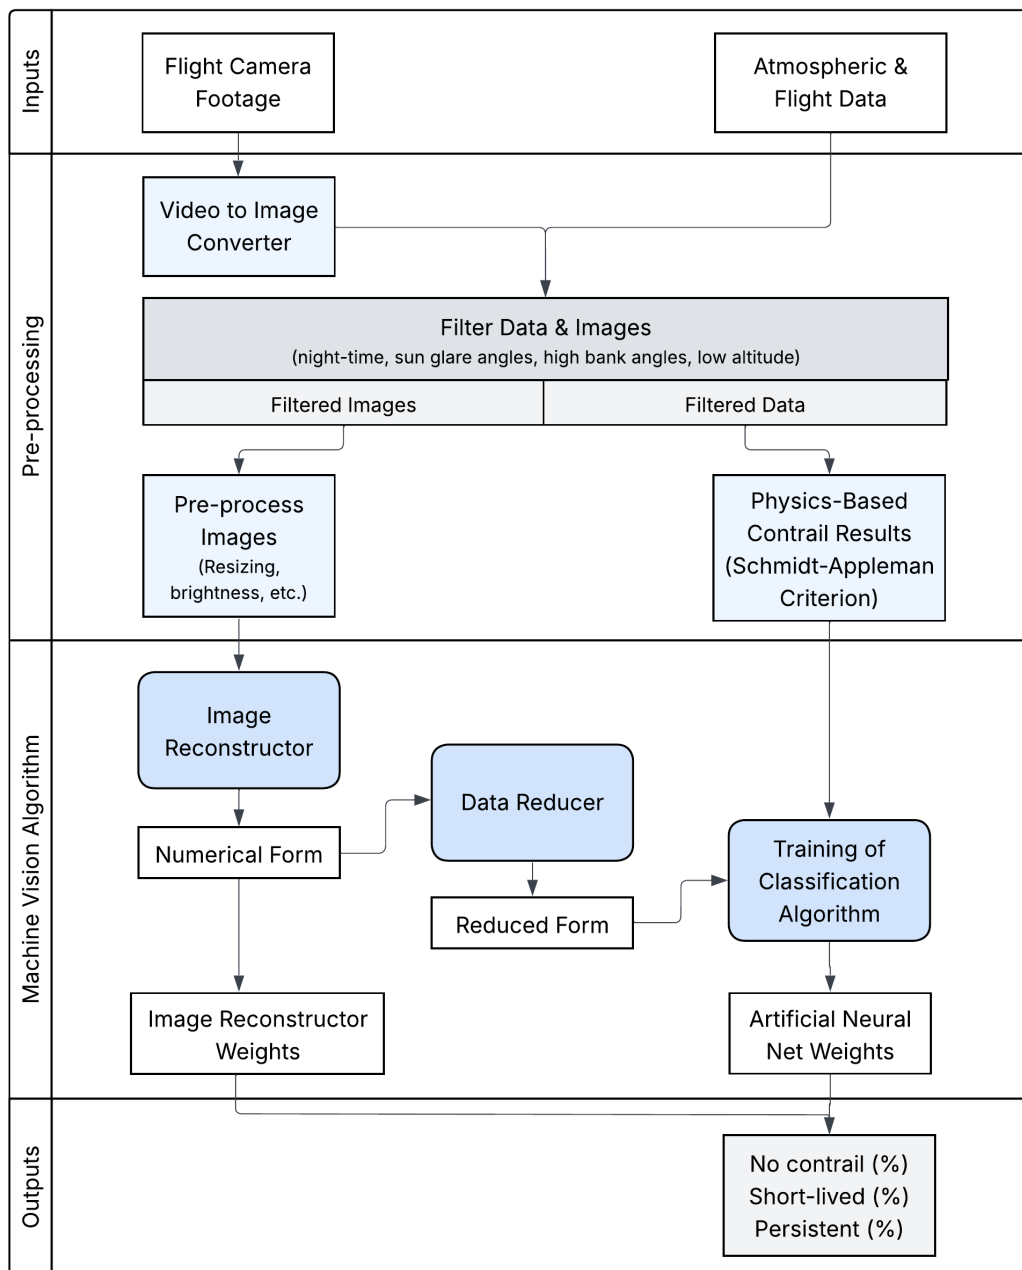
## 135 2 Methodology

This section describes the machine vision framework which ultimately aims to enable onboard contrail detection and prediction of persistence. The methodology builds upon and refines the framework presented in (Papamichou et al., 2026), with special attention given to reconfiguring the preprocessing routine. There are four stages in the methodology, as illustrated in Fig. 1: (a) collection of the inputs used to train the model, (b) filtering and preprocessing of the datasets, (c) training of the machine vision algorithm and finally (d) obtaining the classification results. More specifically, the process starts with camera footage from the aircraft exhaust and a selection of atmospheric and flight data. Then this footage is discretized into image frames to be conveyed through an autoencoder, followed by Principal Component Analysis (PCA), for further dimensionality reduction. The resulting reduced vectors, along with the atmospheric and flight data, are used to train the contrail classification model. The model classifies contrails into three categories: absent, short-lived, and long-lived (persistent).

After the model has been trained, only video footage of the aircraft exhaust is required as input and the operation is simplified to the four steps shown in Fig. 2. The image reconstructor and artificial neural network (ANN) weights refer to the trainable parameters from the autoencoder and ANN, respectively. In the end, the classification results are presented in the form of percentages indicating the likelihood of each contrail category at a given time. A higher percentage in the persistence category indicates a greater expectation that the contrail will persist. This classifier may be leveraged by airlines for operational contrail monitoring, to inform real-time decisions on altitude and route when encountering ISSRs, in order to prevent persistent contrail formation.

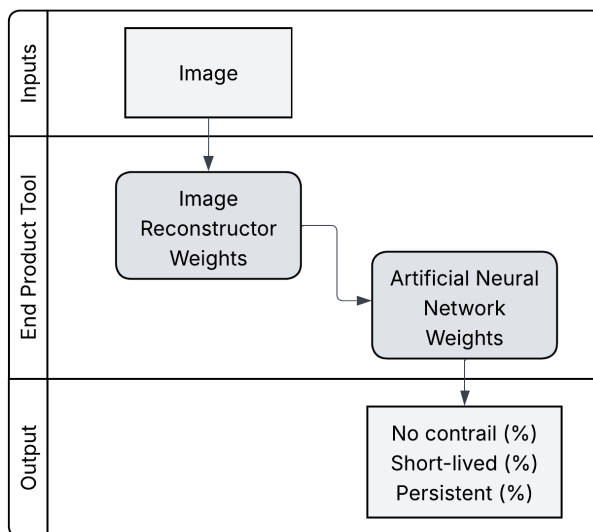
### 2.1 Flight campaigns and input data

This study aims to study the viability of onboard video footage for contrail detection and classification, as well as provide recommendations on the optimal positioning of the camera, which is discussed in Sect. 3. The key input is the video footage of the flight with a clear view of the engine exhaust. The second set of inputs includes the atmospheric and flight path data which correspond to the timeframe of the video footage. The data contain recorded parameters such as time, altitude, latitude, and longitude from flight data, as well as air temperature, air pressure, and relative humidity over ice from atmospheric data.



**Figure 1.** Overall methodology

The data used in this study were collected from two experimental campaigns shared by NASA. During the first campaign, NASA/DLR Multidisciplinary Airborne eXperiments (ND-MAX), the lead aircraft, the Airbus A320-232 D-Advanced Technology Research Aircraft (ATRA) with its IAE V2527-A5 engines, flew over Germany in 2018 collecting emissions data



**Figure 2.** Simplified end-product methodology using the trained autoencoder and ANN, requiring only video footage as input.

followed by NASA’s DC-8 Airborne Science Laboratory chase aircraft. In total, the available data consist of 7 days (approximately 29 hours in total) of flight and atmospheric measurements. The second campaign was a collaboration with Boeing’s ecoDemonstrator aircraft in 2023 over Washington state in the USA. The lead aircraft was the Explorer B737-10 with the CFM LEAP-1B engines followed by NASA’s DC-8. The ecoDemonstrator measurement campaign spanned over 4 days, and yielded approximately 19 hours of flight time.

Both campaigns recorded visual footage from cameras in the nadir, forward, port, and starboard directions. The algorithm requires a view of the contrail behind the engine exhaust and as such either port side or starboard may be utilized due to symmetry. Similar flight and atmospheric variables were measured between campaigns, including, time, altitude, latitude, longitude, air temperature, static pressure, and relative humidity, which were particularly significant.

## 2.2 Preprocessing

The acquisition and preparation of the footage and data require significant attention, since model performance is heavily constrained by the quality of the training data. Three main strategies were employed for the preparation of the flight and atmospheric data: (1) filtering instances with occluded or ambiguous contrails, (2) temporally linking the data to the associated image, and (3) determining the ground truth contrail status labels, based on the atmospheric thermodynamic conditions, to train against. From the ND-MAX dataset, the footage was converted into a set of images taken every 5 seconds, synchronized with the filtered flight and atmospheric data, and resized to ensure uniformity. The frames meeting any of the criteria listed below were removed from the training dataset, initially by filtering through the corresponding flight atmospheric conditions, followed by manual visual inspection of the images:



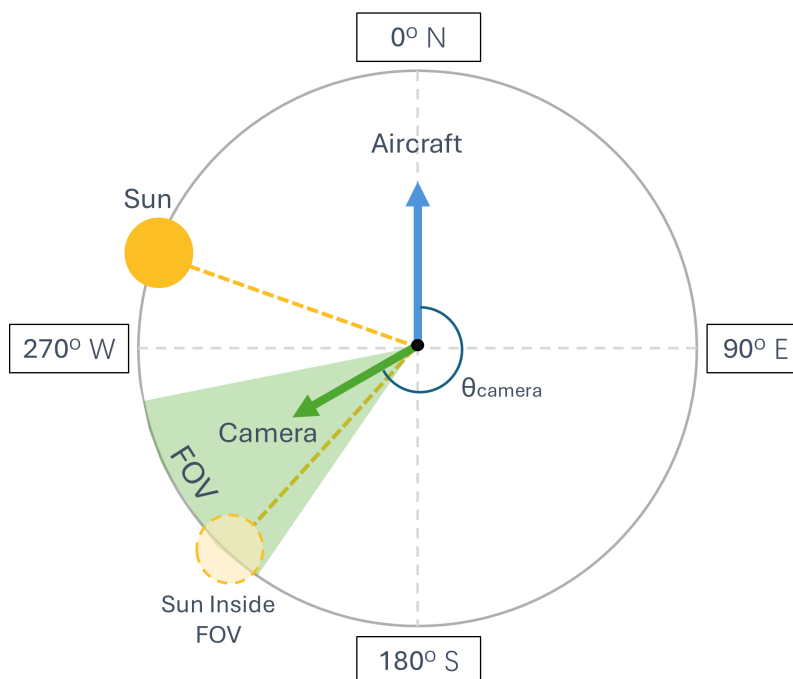
1. frames captured after dusk due to insufficient natural lighting for visual contrail detection,
- 180 2. low-altitude flight segments, below 5000 ft (1524 m), due to occlusion from clouds, and lower likelihood of contrail formation,
3. high bank angles, above  $10^\circ$  due to obstruction of the contrail by the horizontal stabilizer at the tail (only for ND-MAX),
4. sun glare, sun within  $\pm$  certain degrees of the camera's field of view (FOV) for darkening the image and hiding the contrail features.

185 The sun glare filtering required the manual identification and derivation of the relationship between the position of the sun and the occurrence of sun glare, with examples shown in Fig. 3, which caused localized image saturation that obstructed contrail features. The available flight data included the sun azimuth angle and the aircraft's true heading, from which the sun's position relative to the camera's FOV could be geometrically derived, as illustrated in Fig. 4. This enabled the removal of instances with sun interference. The camera heading was assumed to have a symmetric triangular FOV. While the equations  
190 are expected to be consistent among flights, the bounds for the FOV and camera heading angle may change according to the camera and its placement.



**Figure 3.** Camera view of two cases of sun glare. The left image shows an occluded contrail due to sun glare, while the right image retains visibility of the contrail despite the sun glare due to the positioning of the sun.

The actual camera heading is unknown so it was estimated through visual inspection of the images at  $95^\circ$  from the true heading for the ND-MAX measurement data and  $75^\circ$  for the ecoDemonstrator. These values were purposefully selected at an offset from the camera heading, in order to focus on symmetrically capturing the sun glare that occurs when the sun is on the right half of the frame, thus causing a glare in view. Sun glare was determined by taking the angular difference between the sun azimuth and the camera heading. Since the sun glare was largely independent of the vertical axis in the footage from the  
195



**Figure 4.** Directions and angles used for sun glare filtering (top view).

two campaigns, the problem was solved in a two-dimensional space. The filtering conditions were modified to adapt to the uniqueness of the campaign footage. A summary of the differences between the ND-MAX and ecoDemonstrator campaigns are shown in Table 1.

**Table 1.** Filtering conditions for the 2018 ND-MAX and 2023 ecoDemonstrator campaigns

ND-MAX	ecoDemonstrator
- Filter night-time footage	- No dark night-time footage
- Filter below 5000 ft altitude	- Filter cases below 5000 ft altitude
- Filter over 10° bank angle turns	- No contrail occlusion due to tail at high bank angles
- Filter sun within $\pm 45^\circ$ FOV	- Filter sun within $\pm 25^\circ$ FOV
- Camera heading at $95^\circ$	- Camera heading at $75^\circ$

200 The amount of filtering necessary may be decreased by optimizing camera placement and heading. An initial list of suggestions is presented in Sect. 4 with the adjustment of camera placements. Upon completion of the filtering process, the images were synchronized to the filtered atmospheric data in terms of time. As mentioned earlier, during the video to image extraction and conversion, images were stored at a fixed 5-second interval, and the atmospheric data were resampled to be synchronized



with this time interval. Additionally, the atmospheric data were filtered to exclude the aforementioned low visibility conditions,  
205 along with the corresponding images. The synchronization relied on the timestamp of each time frame, which was extracted  
using pytesseract, a Python library wrapper to the Tesseract Optical Character Recognition (OCR) (Smith, 2007).

Editing the images prior to passing them to the autoencoder ensured consistent inputs for the autoencoder architecture that  
is further explained later in this section. This is the final preprocessing step for the images and was achieved by resizing them  
all to a uniform size.

210 The atmospheric data were used in the SAc calculations to determine the expected contrail status. As mentioned in Sect. 1,  
the SAc is a widely accepted thermodynamic criterion for the detection of contrail formation. Following the process described  
in (Dischl et al., 2022), the criterion was evaluated at each timestep along the flight. The emission index (EI) of water vapor,  
the overall pressure ratio of the aircraft engines, and the heat of combustion metrics were assumed to be constant. The air  
temperature and air pressure throughout the flight were extracted from the atmospheric dataset to test against the SAc. Persistent  
215 contrails require a relative humidity over ice of the ambient air over 100 %. Therefore, a relative humidity over ice above the  
threshold signified long-lived contrail formation. A similar approach is followed by pycontrails's CoCiP model, a well-known  
Python model that combines the SAc with the relative humidity over ice threshold.

From this analysis, each image frame is assigned a 0 if a contrail is absent, 1 if a contrail forms and is short-lived, or 2 if it  
forms and is persistent. The resulting matrix, including time and contrail status, is entered into the final classification algorithm  
220 for training.

### 2.3 Machine vision algorithm

A fully convolutional residual autoencoder is developed to preserve the critical characteristics of the image frames before they  
proceed to the classification model with the SAc results. The frames before resizing consist of 720 x 720 pixels and RGB  
channels leading to a 1 555 200 element vector with primarily noisy or non-relevant patterns. Therefore, the autoencoder is  
225 needed to identify only the main patterns and reduce the dimensionality of the images to quickly and accurately prepare these  
series of images. The encoder consists of five convolutional blocks of kernel size 4, stride 2, and padding 1 with channel  
progression from 3 to 512, to represent the substantial visual variability present in cloud scenes. Two residual blocks follow,  
operating at the bottleneck to preserve fine-scale structure, which is critical in modeling correlated atmospheric textures rather  
than discrete objects. The decoder mirrors the encoder using bilinear upsampling followed by convolution, with U-Net-style  
230 skip connections at each scale. A sigmoid activation is used in the final layer to reconstruct the images. Latent vectors are  
extracted producing a 512-dimensional embedding per image frame with adaptive average pooling used in the bottleneck  
representation.

Training uses AdamW with gradient clipping, learning-rate reduction on plateau, and early stopping to allow the latent space  
to organize itself according to the prominent visual and structural features. A composite loss function combining mean squared  
235 error with 55%, L1 loss with 15%, and a structural similarity index metric, SSIM, with 30% components is used to balance  
fine detail preservation and perceptual and structural fidelity. Training ran for up to 50 epochs with batch size 8. After training,  
latent vectors are standardized and analyzed using principal component analysis, PCA. PCA is employed to identify orthogonal



and therefore uncorrelated features within the purely numerical latent space (Chatterjee and Choudhury, 2025). PCA outputs lower-dimensional features that preserve the variance in the structure with correct tuning, and thus the outputs are directly usable by the ANN. Based on trials performed, the top 10 components are chosen, meaning that 96% of the dataset variance is included. All latent embeddings, model weights, and derived statistics are saved to enable exact replication and downstream analysis.

Finally, the classification model is designed to categorize contrail status into three discrete classes. The 10-dimensional input feature vectors from the PCA stage are standardized to zero mean and unit variance. A weighting system is added to address class imbalance, since not all three classes appear in the same number of frames. Those are calculated for each class  $c$  using the formula:

$$w_c = \left( \frac{1}{N_c} \right)^\alpha \quad (1)$$

where  $N_c$  is the class count and  $\alpha$  is a scaling hyperparameter, followed by normalization to ensure the weights sum to the number of classes.

The neural network uses the SAc results to train to classify the reduced images according to contrail status category. The model's architecture consists of a two-layered feature-processing block and a two-layered classification block with dropout layers to mitigate overfitting and concludes in a 3-neuron output layer with a softmax activation. The aim is to balance classification performance and computational time for each flight campaign. A k-fold cross-validation strategy is implemented with the Adam optimizer for training. The final model weights are selected from the fold with the highest validation accuracy. The results presented in Sect. 3 compare the actual versus predicted classes in every contrail category for each image, meaning each image is associated with a vector of three percentages of confidence for absence, short-lived, and persistent.

### 3 Results

As mentioned in the previous sections, the camera footage and corresponding flight and atmospheric data analyzed in this study were collected from the ND-MAX and ecoDemonstrator flight tests. The port or starboard side facing footage from the aircraft was preferred for observing contrail formations at the engine wake. In this section, the results from each part of the methodology illustrated in Fig. 1 are presented and explained.

Both the 2018 ND-MAX and the 2023 ecoDemonstrator campaigns used NASA's DC-8 as the chase aircraft. The measurements obtained from the DC-8 chase aircraft for each flight campaign, consisting of air temperature, static pressure, and relative humidity over ice, were used as inputs to the SAc and ice supersaturation criterion models along the flight paths to determine whether they meet the conditions for contrail formation and persistence, respectively. Certain parameters were assumed to be constant for both the A320 and B737-10 aircraft and conventional jet fuel values from Dischl et al. (2022) were used. These constant parameters include the emissions index of water vapor used, the heat of combustion, and the overall propulsion efficiency. Table 2 contains a summary of all constants used for the SAc evaluation.

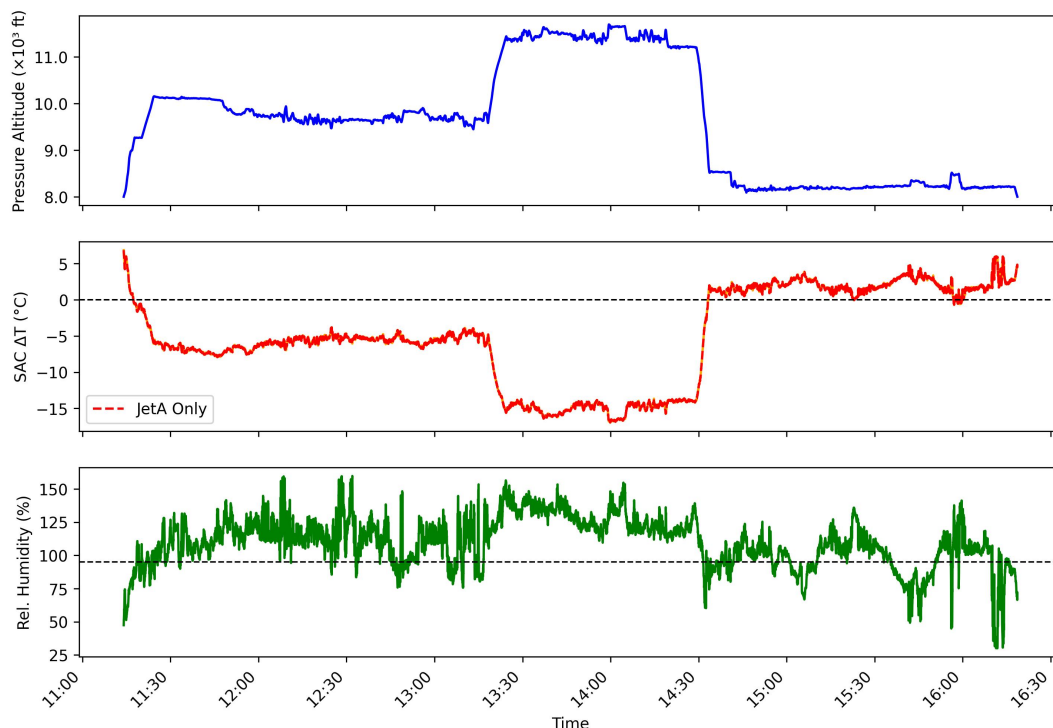


Both campaigns burned Jet-A fuels and SAF (Bräuer et al., 2021; Schripp et al., 2022; Gallo et al., 2024) alternatively during  
270 flights. The model was first tested assuming conventional Jet-A heat of combustion and EI of water vapor ( $EI_{H_2O}$ ) throughout  
the entire duration of the flights. Further testing was performed by implementing the EcoDemonstrator's fuel schedule in the  
SAC step with the appropriate values for  $EI_{H_2O}$  and combustion heat for conventional Jet-A and 100% HEFA SPK Paraffinic  
SAF (Gallo et al., 2024). The ANN classification results from both tests were indistinguishable for the most part, since as  
275 indicated in Table 2, the fuel-dependent parameters differ by as small margin. The evaluation of those parameters followed  
the practice shown in CoCiP. Unlike CoCiP, this model utilizes the SAC to determine if contrail formation is possible, but it  
does not account for the non-volatile particles that SAF blends significantly reduce, allowing them to produce less contrails.  
Therefore, to demonstrate the implementation of this methodology, the following results were created with solely Jet-A values  
throughout all flights.

**Table 2.** Parameters used in the Schmidt-Appleman criterion evaluation. Values that differ between Jet A and 100 % SAF blend are shown separately; all other parameters are common to both fuel types.

Parameter	Unit	Jet A	100 % SAF
<i>Fuel-dependent parameters</i>			
Emission index of water vapor	kg kg <sup>-1</sup>	1.23	1.36
Heat of combustion	MJ kg <sup>-1</sup>	43.130	43.131
<i>Common parameters</i>			
Propulsion efficiency	–		0.33
Molar mass ratio (H <sub>2</sub> O/air)	–		0.622
Specific heat capacity of air	J kg <sup>-1</sup> K <sup>-1</sup>		1004

Figure 5 illustrates the conditions during the ND-MAX flight on 17 January 2018. For the pressure altitude at every time step,  
280 the SAC  $\Delta T$  plots the temperature difference between the ambient air and the threshold temperature for contrail formation. A  
negative  $\Delta T$  indicates that the ambient temperature is below the threshold, suggesting conditions that favor contrail formation.  
The relative humidity over ice plot marks times during the flight that exceeded the 95 % threshold and are thus more likely  
to produce persistent contrails. The critical relative humidity for contrail formation was set at 95 %, following (Dischl et al.,  
2022), to allow some margin from the theoretical threshold of 100 %. For ND-MAX, the likelihood of contrail formation  
285 appears above 8000 ft between approximately 11:00 and 14:30 UTC. Persistence is possible at several points across the flight,  
with multiple peaks above the threshold. Relative humidity remains high after 14:30 UTC but the positive SAC  $\Delta T$  deems a  
contrail unlikely. A corresponding sample from the ecoDemonstrator on 21 October 2023 is shown in Fig. 6. The SAC  $\Delta T$  plot  
shows the minimal impact of selecting Jet A instead of the actual fuel schedule for the SAC calculation. As illustrated in Fig. 1,  
the classification neural network trains by leveraging these physics-based results.

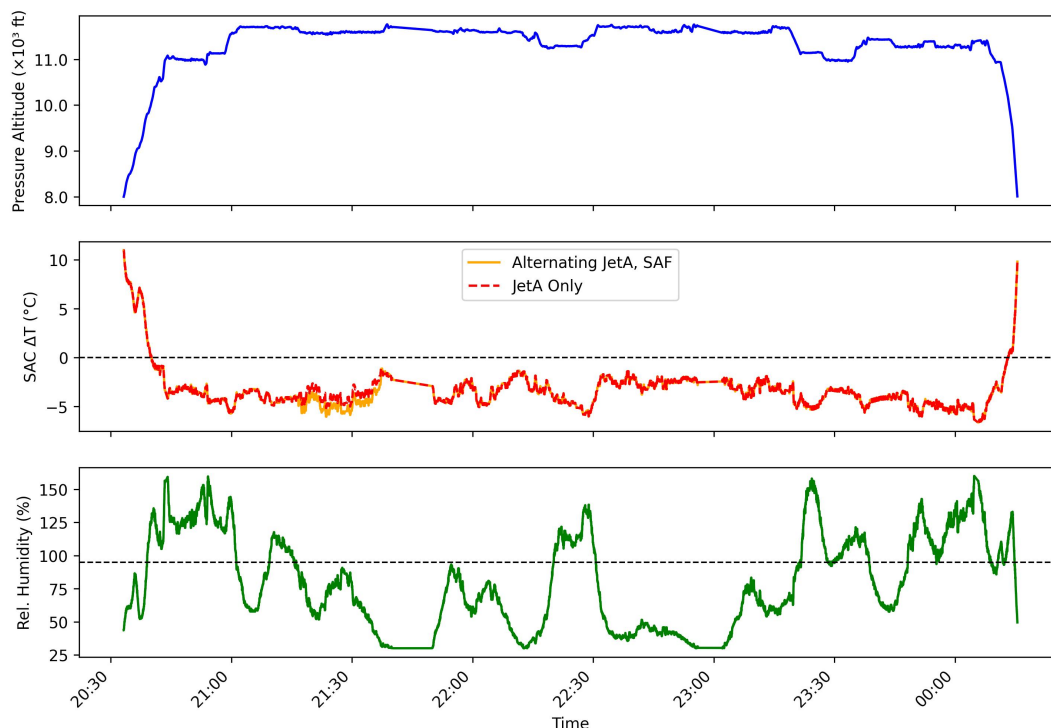


**Figure 5.** Temperature difference between ambient and threshold for SAC (SAC  $\Delta T$ ) (top) and relative humidity over ice (bottom) for the pressure altitudes from ND-MAX on 17 January 2018.

290 Certain conditions were observed in the camera footage that contribute to reduced visibility, decreasing the effectiveness of the model, as detailed in Sect. 2. Filtering was employed to exclude those cases. Those conditions are highly dependent on camera placement and direction, and several conclusions can be deduced to improve them and increase model accuracy.

The filtering process described in Sect. 2.2 reduced the atmospheric data from approximately 131 000 to 39 000 samples for ND-MAX and from 139 000 to 57 000 for the ecoDemonstrator. The video was then discretized and synchronized with the atmospheric data. The footage was converted to images at 5-second and 2-second intervals apart for ND-MAX and ecoDemonstrator, respectively, the latter chosen to be a shorter interval to account for the shorter duration of footage. The result consisted of roughly 17 200 and 24 000 images, respectively. A synchronization tool was created to match and store the atmospheric data corresponding to the remaining time instances of the images. Furthermore, it deleted the images corresponding to the filtered atmospheric data. The synchronization reduced the usable number of images and data to 7 917 for ND-MAX and 16 561 for ecoDemonstrator.

300 The image frames were resized from  $720 \times 720$  to  $320 \times 320$  pixels to reduce the input vector to a size that can be processed by an onboard computer. In this study, image editing was not performed. The autoencoder was trained for 50 epochs with early stopping for both campaigns. The training metrics shown in Table 3 suggest stable convergence with training loss below 0.002,



**Figure 6.** Temperature difference between ambient and threshold for SAC (SAC  $\Delta T$ ) (top) and relative humidity over ice (bottom) for the pressure altitudes from ecoDemonstrator on 21 October 2023.

and high reconstruction quality with PSNR above 37 dB and SSIM close to 1. The resulting latent space was normalized before  
 305 entering the PCA to encourage more pronounced clustering.

**Table 3.** Autoencoder training metrics for the two campaigns.

Campaign	Training Loss	PSNR	SSIM
ND-MAX	0.00112	39.97	0.99998
ecoDemonstrator	0.00113	37.08	0.99910

The classification was performed using an ANN evaluated with a 5-fold cross validation to improve the robustness of the evaluation and reduce the risk of overfitting. The network’s architecture included an input layer that mapped 10-dimensional principal component vectors into a 32-dimensional feature space, followed by a sequence of fully-connected layers with layer normalization, ReLU activation functions and dropout for training stability and regularization. The resulting model is evaluated  
 310 through inference on the dataset to determine accuracy and converge on the best performing model. For the training process, the hyperparameter  $\alpha$  was set to 0.5 to most appropriately address the class imbalance from the available data. Table 4 summarizes



the number of images per class and the weight assigned to them to avoid an imbalanced dataset. The Adam optimizer was used with a learning rate of 0.001 and cross-entropy loss. Training employed 5-fold cross-validation with early stopping, halting if validation loss did not improve by more than 0.001 over 30 consecutive epochs. Out-of-fold predictions were collected so that each of the 14,470 or 16,561 campaign samples was classified exactly once, during its respective held-out fold. The final per-class metrics and confusion matrix were computed over the complete set of out-of-fold predictions across all samples.

**Table 4.** Image count per class and the weighting assigned to them to avoid class bias due to data imbalance.

<b>ND-MAX</b>			
<i>Metric</i>	<i>Absent</i>	<i>Short-lived</i>	<i>Persistent</i>
Class Count	3453	5778	5239
Class Weight	1.1605	0.8972	0.9422
<b>ecoDemonstrator</b>			
Class Count	4229	6827	5505
Class Weight	1.1263	0.8865	0.9872

Table 5 reports per-class precision, recall, and F1-score derived from out-of-fold predictions across all five folds, such that each sample was evaluated exactly once on unseen data. The model performs consistently across all three contrail classes. F1-scores range from 0.827 to 0.932. The absent class achieves recall values above 0.90, indicating the model does not disproportionately misclassify underrepresented samples. The weighted-average F1 of 0.8671 and 0.9099 for ND-MAX and ecoDemonstrator, respectively, accounts for class frequency.

**Table 5.** Classification metrics by contrail formation class

<b>Class</b>	<b>ND-MAX</b>			<b>ecoDemonstrator</b>		
	<b>Precision</b>	<b>Recall</b>	<b>F1</b>	<b>Precision</b>	<b>Recall</b>	<b>F1</b>
Absent	0.9256	0.9401	0.9328	0.9021	0.9045	0.9033
Short-lived	0.8571	0.8712	0.8641	0.9006	0.8950	0.8978
Persistent	0.8391	0.8152	0.8270	0.9274	0.9328	0.9301
Weighted avg	0.8669	0.8674	0.8671	0.9099	0.9100	0.9099

Confusion matrices were constructed to gauge the model’s accuracy. Table 6 shows the confusion matrix from the ND-MAX campaign, while Table 7 shows the matrix from the ecoDemonstrator results. The percentages represent the ratio of data in a given actual state that the model predicts. For example, the top right cell denotes the percentage of data points where the contrail is absent but is perceived by the algorithm to be persistent. The aim is to approach 100% in the diagonal starting at top left and ending in bottom right with 0% elsewhere.



**Table 6.** Confusion matrix for the machine vision algorithm applied to the ND-MAX campaign.

		Predicted		
		Absence	Short-lived	Persistent
Actual	Absence	94.0 %	1.9 %	4.1 %
	Short-lived	1.1 %	87.1 %	11.8 %
	Persistent	3.7 %	14.7 %	81.5 %

**Table 7.** Confusion matrix for the machine vision algorithm applied to the ecoDemonstrator campaign.

		Predicted		
		Absence	Short-lived	Persistent
Actual	Absence	90.4 %	8.3 %	1.2 %
	Short-lived	5.4 %	89.5 %	5.1 %
	Persistent	0.9 %	5.8 %	93.3 %

#### 4 Discussion

The results show a strong potential for future deployment of the classifier onboard aircraft, along with cameras that capture contrails forming. Through a comparison of the results from the two flight measurement campaigns, the classifier appears to predict contrail absence cases with slightly higher accuracy for the ND-MAX campaign (94.0 % versus 90.4 % for the ecoDemonstrator dataset) and relatively similar accuracy for short-lived contrails (87.1 % for ND-MAX versus 89.5 % for the ecoDemonstrator). The slightly lower accuracy in absence for the ecoDemonstrator campaign, despite a higher number of images with absence, may be partially attributed to the camera orientation, which captured a short segment of the exhaust plume close to the engine. Since the distance from the engine at which contrails are formed varies, some contrails could have been out of the camera’s FOV. While there is still room for improvement, these results show a notable improvement over the ND-MAX prototype results that were presented in (Papamichou et al., 2026), which showed 81.3 % and 77.5 % for short-lived and persistent contrails, respectively. Results showed improvement with adjustments to the hyperparameters and filtering precision. The classification rates of both measurement campaigns presented in this work exceed the 50 % accuracy that has been achieved by other current machine learning approaches for contrail identification from satellite imagery of clouds, as documented by the (International Air Transport Association, 2024) report.

This study improved on previous work (Papamichou et al., 2026) in a robust and automated framework. The similarities between the ND-MAX and ecoDemonstrator results and footage, indicate the suitability of the footage and camera placement. Their slight differences suggest an impact of the camera placement and footage quality on the model’s accuracy. Optimizing

camera orientations in the future will be a significant factor in improving the training process and the classification accuracy  
345 of the model in future deployments, particularly for persistence.



**Figure 7.** Comparison of the ND-MAX and ecoDemonstrator footage. (Top Left) ND-MAX campaign footage has tail obstruction in the view of the contrail, as well as a visible large portion of the aircraft fuselage and wing. (Top Right) ecoDemonstrator footage is not occupied by the aircraft body and the tail never covers the contrail stream. (Bottom Left) ND-MAX campaign footage is typically lower brightness and has a wider FOV, increasing the visible length of the contrail stream. (Bottom Right) ecoDemonstrator footage has high brightness and shorter view of the contrail.

Testing the algorithm on the ND-MAX and ecoDemonstrator footage has revealed certain optimal visual conditions. As shown in the images from the two campaigns in Fig. 7, there are notable differences. In the ND-MAX campaign video, the aircraft's tail obstructs a portion of the contrail at high roll angles due to its camera placement. On the other hand, the camera location used in the ecoDemonstrator campaign is lower and angled outward, thus avoiding the aforementioned issue.  
350 Additionally, a significant part of the fuselage and wing occupies the captured ND-MAX frame, unlike in the ecoDemonstrator footage, where the outward orientation of the camera minimizes intrusion from the airframe. However, the ecoDemonstrator's camera angle shows a significantly shorter segment of the contrail stream, whereas the ND-MAX footage, extends to several kilometers behind the aircraft in favorable conditions.



The results for persistence are better for the ecoDemonstrator footage compared to the ND-MAX. Therefore, it may be  
355 advisable to locate camera placements that minimize the airframe components' visibility, since they further complicate a  
challenging task. At the same time, the absence has higher accuracy for the ND-MAX, thus, maximizing the length of the  
contrail, similar to ND-MAX, could be beneficial. The length of a contrail renders it more distinguishable among natural  
clouds. Additionally, placements should be considered that avoid being at a similar height to the horizontal stabilizer. Either  
a higher or lower placement should prevent the tail from hiding the contrail at high roll angles. For sun glare, filtering might  
360 be reduced by placing the camera high and pointing it downwards or tinting the top half of the FOV. An additional camera on  
the opposite side may help capture the cases where sun glare is too obtrusive on one side's footage. Lastly, a separate study  
may be conducted for the use of chase aircraft forward facing footage with this paper's methodology to investigate the effect  
of showing the model a larger segment of the contrail on the persistence prediction.

## 5 Conclusions

365 This work presented a methodology that uses machine vision algorithms to enable real-time onboard contrail detection and  
persistence prediction, using footage from two airborne flight campaigns to train the respective models. This methodology  
combines an autoencoder used for feature extraction, PCA for additional dimensionality reduction and an ANN for classifica-  
tion of contrail status from visual footage into three categories: absent, short-lived and persistent. This approach was developed  
and evaluated using publicly available data from two NASA flight campaigns: the 2018 NASA-DLR ECLIF II/ND-MAX over  
370 Germany and the 2023 NASA-Boeing ecoDemonstrator over Washington state.

Two classifiers were trained, one for each flight campaign, with promising results for both campaigns. For the ND-MAX  
dataset, accuracies of 94.0 %, 87.1 % and 81.5 % were obtained for contrail absence, short-lived and persistent categories,  
respectively. For the ecoDemonstrator, those values were 90.4 %, 89.5 % and 93.3 %, respectively. These results show that  
implementing onboard cameras with this tool may be beneficial for contrail tracking and persistence prediction. The main  
375 reason for misclassification cases was the confusion between the short-lived and persistent categories, which can be attributed  
to the continuous evolution of the contrail between those two contrail states.

A key finding of the study was the importance of camera placement on classification accuracy. The results suggest that future  
camera installations should aim to maximize the visible contrail length while minimizing the presence of aircraft components in  
the frame. Additionally, the positioning of cameras relative to the horizontal stabilizer should be considered to avoid occlusion  
380 at high roll angles. Reducing instances of sun glare should be attempted depending on placement.

The approach described in this work offers several practical advantages for fleet-wide deployment. It simply requires a  
camera with a clear view of the engine exhaust as additional instrumentation, it can allow for real time operations, and it is  
more computationally lightweight compared to the use of physics-based simulation tools or satellite-based detection systems  
that rely on multiple external data sources. As such, it may be leveraged as a complementary tool to support contrail monitoring  
385 and mitigation strategies.



There are several directions identified for future work. First, the classification accuracy can be evaluated for potential improvement by optimizing camera placement for future airborne measurement campaigns, based on the recommendations presented in Sect. 4, as well as through the further optimization of hyperparameters per flight campaign dataset. Second, the integration of this onboard detection system with atmospheric data and flight planning tools could enable the real-time avoidance of ISSRs, thus preventing the formation of persistent contrails. Third, the methodology could be extended to forward-facing chase aircraft footage to investigate the effect of a larger visible contrail segment on persistence prediction.

*Code and data availability.* The source code to reproduce the methodology described in this paper is available at: <https://doi.org/10.5281/zenodo.20485244>. For ease of reproducibility, the datasets can also be found here in the expected format of the source code: <https://doi.org/10.5281/zenodo.20422913>. The dataset consist of atmospheric measurements and onboard camera footage of contrail formation from the ECLIF II/ND-MAX and ecoDemonstrator flight campaigns found publicly in the NASA archives below. The video footage for the 2018 NASA - DLR ECLIF II/ND-MAX and 2023 NASA - Boeing ecoDemonstrator flight measurement campaigns, as well as the atmospheric measurement data for the 2023 ecoDemonstrator campaign are publicly available and can be found through the NASA Airborne Science Program Online Video Archive (<https://asp-archive.arc.nasa.gov/>), as part of the NASA Airborne Science Program. The atmospheric measurements and aircraft data for the ND-MAX campaign can be found through the NASA Aeronautics Fields Projects, Public Projects List (<https://science-data.larc.nasa.gov/aero-fp/projects/>), as part of the NASA Aeronautics Mission Directorate.

*Author contributions.* AP: methodology, software, investigation, visualization, writing - original draft, writing - review & editing, EK: conceptualization, methodology, investigation, supervision, writing -original draft, writing -review & editing, RM: investigation, resources, writing -review & editing, HP: conceptualization, writing - review & editing, DM: writing - review & editing.

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

*Acknowledgements.* The authors thank Harry Shrager, Petros Famellos, and Lydia M. Szuwalski for their contributions to our codebase and Ryan Bennett from NASA's Bay Area Environmental Research Institute (BAER) for providing the video and atmospheric data.



## References

- Akhtar Martínez, C., Eastham, S. D., and Jarrett, J. P.: Zero-dimensional contrail models could underpredict lifetime optical depth, *Atmospheric Chemistry and Physics*, 25, 12 875–12 891, <https://doi.org/10.5194/acp-25-12875-2025>, 2025.
- 410 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.
- Bräuer, T., Voigt, C., Sauer, D., Kaufmann, S., Hahn, V., Scheibe, M., Schlager, H., Diskin, G. S., Nowak, J. B., DiGangi, J. P., Huber, F., Moore, R. H., and Anderson, B. E.: Airborne Measurements of Contrail Ice Properties—Dependence on Temperature and Humidity, *Geophysical Research Letters*, 48, e2020GL092 166, <https://doi.org/https://doi.org/10.1029/2020GL092166>, 2021.
- 415 Chatterjee, S. and Choudhury, S. J.: Exploring structural components in autoencoder-based data clustering, *Engineering Applications of Artificial Intelligence*, 140, 109 562, <https://doi.org/10.1016/j.engappai.2024.109562>, 2025.
- Dischl, R., Kaufmann, S., and Voigt, C.: Regional and Seasonal Dependence of the Potential Contrail Cover and the Potential Contrail Cirrus Cover over Europe, *Aerospace*, 9, <https://doi.org/10.3390/aerospace9090485>, 2022.
- European Commission, Directorate-General for Climate Action: New monitoring rules agreed for the EU  
420 ETS, including non-CO<sub>2</sub> emissions from the aviation sector, [https://climate.ec.europa.eu/news-your-voice/news/new-monitoring-rules-agreed-eu-ets-including-non-co2-emissions-aviation-sector-2024-08-30\\_en](https://climate.ec.europa.eu/news-your-voice/news/new-monitoring-rules-agreed-eu-ets-including-non-co2-emissions-aviation-sector-2024-08-30_en), last access: 30 March 2026, 2024.
- Freudenthaler, V., Homburg, F., and Jäger, H.: Contrail observations by ground-based scanning lidar: Cross-sectional growth, *Geophysical Research Letters*, 22, 3501–3504, <https://doi.org/https://doi.org/10.1029/95GL03549>, 1995.
- 425 Fritz, T. M., Eastham, S. D., Speth, R. L., and Barrett, S. R. H.: The role of plume-scale processes in long-term impacts of aircraft emissions, *Atmospheric Chemistry and Physics*, 20, 5697–5727, <https://doi.org/10.5194/acp-20-5697-2020>, 2020.
- Gallo, F., Anderson, B. E., Baughcum, S., Brown, M. D., Griffin, W. M., Jordan, C., Shook, M. A., Wiggins, E. B., Winstead, E. L., Ziemba, L. D., and Moore, R. H.: Aircraft Engine Particle Emissions during the 2023 NASA-Boeing ecoDemonstrator Field Measurements, in: AAAR 42nd Annual Conference, Albuquerque, NM, USA, <https://ntrs.nasa.gov/citations/20240013052>, last access: 24 April 2026, 2024.
- 430 Geraedts, S., Brand, E., Dean, T. R., Eastham, S., Elkin, C., Engberg, Z., Hager, U., Langmore, I., McCloskey, K., Yue-Hei Ng, J., Platt, J. C., Sankar, T., Sarna, A., Shapiro, M., and Goyal, N.: A scalable system to measure contrail formation on a per-flight basis, *Environmental Research Communications*, 6, 015 008, <https://doi.org/10.1088/2515-7620/ad11ab>, 2024.
- Google Research: Project Contrails, <https://sites.research.google/contrails/>, last access: 24 April 2026, 2023.
- Hofer, S., Gierens, K., and Rohs, S.: How well can persistent contrails be predicted? An update, *Atmospheric Chemistry and Physics*, 24,  
435 7911–7925, <https://doi.org/10.5194/acp-24-7911-2024>, 2024.
- Hoffman, J. P., Rahmes, T. F., Wimmers, A. J., and Feltz, W. F.: The Application of a Convolutional Neural Network for the Detection of Contrails in Satellite Imagery, *Remote Sensing*, 15, 2854, <https://doi.org/10.3390/rs15112854>, 2023.
- International Air Transport Association: Aviation contrails and their climate effect: Tackling uncertainties and enabling solutions, <https://www.iata.org/contentassets/726b8a2559ad48fe9decb6f2534549a6/aviation-contrails-climate-impact-report.pdf>, last access: 30 March  
440 2026, 2024.
- Kärcher, B.: Formation and radiative forcing of contrail cirrus, *Nature Communications*, 9, 1824, <https://doi.org/10.1038/s41467-018-04068-0>, 2018.



- Lee, D., Fahey, D., Skowron, A., Allen, M., Burkhardt, U., Chen, Q., Doherty, S., Freeman, S., Forster, P., Fuglestvedt, J., Gettelman, A., De León, R., Lim, L., Lund, M., Millar, R., Owen, B., Penner, J., Pitari, G., Prather, M., Sausen, R., and Wilcox, L.:  
445 The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018, *Atmospheric Environment*, 244, 117 834, <https://doi.org/10.1016/j.atmosenv.2020.117834>, 2021.
- Low, J., Teoh, R., Ponsobny, J., Gryspeerdt, E., Shapiro, M., and Stettler, M. E. J.: Ground-based contrail observations: comparisons with reanalysis weather data and contrail model simulations, *Atmospheric Measurement Techniques*, 18, 37–56, <https://doi.org/10.5194/amt-18-37-2025>, 2025.
- 450 Mandija, F., Keckhut, P., Alraddawi, D., Irbah, A., Sarkissian, A., Khaykin, S., Peyrin, F., and Baray, J.-L.: Nighttime Contrail Characterization from Multisource Lidar and Meteorological Observations, *Remote Sensing*, 18, <https://doi.org/10.3390/rs18020210>, 2026.
- Meijer, V. R., Eastham, S. D., Waitz, I. A., and Barrett, S. R. H.: Contrail altitude estimation using GOES-16 ABI data and deep learning, *Atmospheric Measurement Techniques*, 17, 6145–6162, <https://doi.org/10.5194/amt-17-6145-2024>, 2024.
- NASA Aeronautics Mission Directorate: NASA Aeronautics Fields Projects, Public Projects List, <https://science-data.larc.nasa.gov/aero-fp/projects/>, last access: 24 April 2026.
- 455 NASA Airborne Science Program: NASA Airborne Science Program Online Video Archive, <https://asp-archive.arc.nasa.gov/>, last access: 24 April 2026.
- National Academies of Sciences, Engineering, and Medicine: Developing a Research Agenda on Contrails and Their Climate Impacts, The National Academies Press, Washington, DC, <https://doi.org/10.17226/29073>, 2025.
- 460 Ng, J. Y.-H., McCloskey, K., Cui, J., Meijer, V. R., Brand, E., Sarna, A., Goyal, N., Arsdale, C. V., and Geraedts, S.: OpenContrails: Benchmarking Contrail Detection on GOES-16 ABI, <https://doi.org/10.48550/arXiv.2304.02122>, 2023.
- Papamichou, A. K., Famellos, P., Szuwalski, L. M., Kallou, E., and Mavris, D.: Detection of Contrail Formation Using Machine Vision, in: AIAA SCITECH 2026 Forum, <https://doi.org/10.2514/6.2026-0425>, 2026.
- Riggi-Carollo, E., Dubot, T., Sarrat, C., and Bedouet, J.: AI-Driven Identification of Contrail Sources: Integrating Satellite Observation and  
465 Air Traffic Data, *Journal of Open Aviation Science*, 1, <https://doi.org/10.59490/joas.2023.7209>, 2023.
- 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, *Atmospheric Measurement Techniques*, 18, 3495–3532, <https://doi.org/10.5194/amt-18-3495-2025>, 2025.
- Schripp, T., Anderson, B. E., Bauder, U., Rauch, B., Corbin, J. C., Smallwood, G. J., Lobo, P., Crosbie, E. C., Shook, M. A.,  
470 Miake-Lye, R. C., Yu, Z., Freedman, A., Whitefield, P. D., Robinson, C. E., Achterberg, S. L., Köhler, M., Oßwald, P., Grein, T., Sauer, D., Voigt, C., Schlager, H., and LeClercq, P.: Aircraft engine particulate matter emissions from sustainable aviation fuels: Results from ground-based measurements during the NASA/DLR campaign ECLIF2/ND-MAX, *Fuel*, 325, 124 764, <https://doi.org/https://doi.org/10.1016/j.fuel.2022.124764>, 2022.
- Schumann, U.: On conditions for contrail formation from aircraft exhausts, *Meteorologische Zeitschrift*, 5, 4–23, <https://doi.org/10.1127/metz/5/1996/4>, 1996.
- 475 Schumann, U.: A contrail cirrus prediction model, *Geoscientific Model Development*, 5, 543–580, <https://doi.org/10.5194/gmd-5-543-2012>, 2012.
- 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, <https://doi.org/10.5194/amt-6-3597-2013>, 2013.



- 480 Schumann, U., Baumann, R., Baumgardner, D., Bedka, S. T., Duda, D. P., Freudenthaler, V., Gayet, J.-F., Heymsfield, A. J., Minnis, P.,  
Quante, M., Raschke, E., Schlager, H., Vázquez-Navarro, M., Voigt, C., and Wang, Z.: Properties of individual contrails: a compilation of  
observations and some comparisons, *Atmospheric Chemistry and Physics*, 17, 403–438, <https://doi.org/10.5194/acp-17-403-2017>, 2017.
- Singh, D. K., Sanyal, S., and Wuebbles, D. J.: Understanding the role of contrails and contrail cirrus in climate change: a global perspective,  
*Atmospheric Chemistry and Physics*, 24, 9219–9262, <https://doi.org/10.5194/acp-24-9219-2024>, 2024.
- 485 Smith, R.: An Overview of the Tesseract OCR Engine, in: Ninth International Conference on Document Analysis and Recognition (ICDAR)  
2007, pp. 629–633, IEEE Computer Society, <https://doi.org/10.1109/ICDAR.2007.4376991>, 2007.
- Soleimanpour, M., Seelig, T., Groß, S., and Tesche, M.: Detection of embedded contrails in airborne lidar measurements, *Atmospheric  
Measurement Techniques*, 19, 3291–3308, <https://doi.org/10.5194/amt-19-3291-2026>, 2026.
- Teoh, R., Schumann, U., Gryspeerdt, E., Shapiro, M., Molloy, J., Koudis, G., Voigt, C., and Stettler, M. E. J.: Aviation contrail climate effects  
490 in the North Atlantic from 2016 to 2021, *Atmospheric Chemistry and Physics*, 22, 10919–10935, <https://doi.org/10.5194/acp-22-10919-2022>, 2022.
- Voigt, C., Kleine, J., Sauer, D., Moore, R. H., Bräuer, T., Le Clercq, P., Kaufmann, S., Scheibe, M., Jurkat-Witschas, T., Aigner, M., Bauder,  
U., Boose, Y., Borrmann, S., Crosbie, E., Diskin, G. S., DiGangi, J., Hahn, V., Heckl, C., Huber, F., Nowak, J. B., Rapp, M., Rauch, B.,  
Robinson, C., Schripp, T., Shook, M., Winstead, E., Ziemba, L., Schlager, H., and Anderson, B. E.: Cleaner burning aviation fuels can  
495 reduce contrail cloudiness, *Communications Earth & Environment*, 2, 114, <https://doi.org/10.1038/s43247-021-00174-y>, 2021.