

Quantifying the Observability of Contrails in Satellite Images

Marianna Piperigou Grammatika

Quantifying the Observability of Contrails in Satellite Images

by

Marianna Piperigou Grammatika

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday July 1, 2026 at 9:00.
Faculty of Aerospace Engineering TU Delft, Lecture Hall C

Student number: 5236398
Project duration: September, 2026 – June, 2026
Thesis committee: Dr. F. Yin, TU Delft, Chair
Dr. V.R. Meijer, TU Delft, Supervisor
Dr.ir. P. Proesmans, TU Delft, Examiner

Contents

List of Abbreviations	4
List of Figures	6
List of Tables	9
1 Introduction	10
2 Literature Review	12
2.1 Schmidt-Appleman criterion	12
2.2 Contrail Cirrus Prediction model	13
2.3 Satellite-based contrail detection	14
2.4 Ground-based camera contrail detection	15
2.5 Flight to contrail attribution	16
2.6 Research gap	17
3 Data Description	18
3.1 All-Sky Imager	18
3.1.1 Contrail Activity Time Series and Case Selection	19
3.2 Satellite Imagery	19
3.2.1 GOES-East ABI	19
3.2.2 MTG FCI	21
3.3 ERA5 Reanalysis	21
3.4 ADS-B Flight Data	22
3.4.1 Flight Track Advection	22
3.5 VOCASION Description	22
4 Methodology	25
4.1 Contrail Width and Age Estimation	25
4.1.1 All-Sky Imager Setup	25
4.1.2 Fisheye Camera Model	26

4.1.3	Image Rectification	27
4.1.4	Contrail Polygon Annotation	28
4.1.5	Projection to Physical Coordinates	29
4.1.6	Altitude and Age Estimation	30
4.1.7	Width Estimation	31
4.1.8	Width Metrics	32
4.1.9	Sources of Uncertainty	34
4.2	Ambient Condition Estimates	35
4.2.1	Relative Humidity with Respect to Ice	35
4.2.2	Vertical Wind Shear	36
4.2.3	Atmospheric Stability	37
4.3	Visibility in Satellite Imagery	37
4.3.1	Parallax Correction	37
4.3.2	Brightness Temperature Calculation	38
4.3.3	Effective Resolution at the Sites	40
4.3.4	Background Condition Labels	41
5	Results	42
5.1	Dataset Overview	42
5.1.1	Collected Variables	44
5.2	What Determines Observability?	44
5.2.1	Visibility Label	45
5.2.2	Contrail Width	46
5.2.3	Age at First Observation	47
5.2.4	Age-width relationship	49
5.2.5	Brightness Temperature Difference at First Observation	50
5.3	What Affects the Delay?	57
5.3.1	Age and Spreading Rate	57
5.3.2	Ambient Conditions	58
5.3.3	Background Conditions	61
6	Conclusion	63
6.1	Recommendations for future work	64

List of Abbreviations

- ABI** Advanced Baseline Imager. 6, 19, 20
- ACTA** Automatic Contrail Tracking Algorithm. 14
- ADS-B** Automatic Dependent Surveillance-Broadcast. 15
- ARM** Atmospheric Radiation Measurement. 18
- ASI** All Sky Imager. 18–20, 22–25, 42, 46, 58, 64
- BTD** brightness temperature difference. 7, 8, 20, 21, 50–57
- CAB** Cabauw. 18, 63
- CDA** Contrail Detection Algorithm. 14, 15
- CO₂** carbon dioxide. 10
- CoCiP** Contrail Cirrus Prediction. 13
- ECMWF** European Centre for Medium-Range Weather Forecasts. 21
- FAA** Federal Aviation Administration. 59
- FCI** Flexible Combined Imager. 21
- FOV** field-of-view. 19, 28, 42, 46
- HR** high-resolution. 54, 57
- ICAO** International Civil Aviation Organization. 59
- ISSR** ice-supersaturated region. 10, 12
- LEO** Low Earth Orbit. 11, 14
- MRR** minimum rotated rectangle. 32, 36
- MSL** mean sea level. 30
- MTOW** Maximum Take Off Weight. 59
- SAC** Schmidt-Appleman Criterion. 12, 13

SGP Southern Great Plains. 18

SRF Spectral Response Function. 54

UI User Interface. 6, 22, 23

WTC Wake Turbulence Category. 59

List of Figures

2.1	Conditions for contrail formation according to the Schmidt-Appleman criterion [9]	13
3.1	Illustration of the event-definition rule. Each rectangle represents one 10-minute bin. Blue bins contain at least one contrail detection; gray bins contain none. A sequence of consecutive blue bins with no intervening gap forms a single event (Event 1, Event 2, Event 3). A single grey bin is sufficient to terminate the current event and start a new one.	19
3.2	GOES-East Advanced Baseline Imager (ABI) Mode 6 scanning: spatial coverage and timeline [40]	20
3.3	VOCASION User Interface (UI)	23
3.4	VOCASION polygon example	23
4.1	Fisheye camera projection geometry. Taken from Kannala and Brandt [47].	26
4.2	Pinhole camera model. A real-world point $P(X, Y, Z)$ is projected through the camera onto the image plane at pixel $p'(u, v)$, at radial distance r from the principal point (u_0, v_0) and viewing angle θ_p from the optical axis.	27
4.3	Raw fisheye image (left) and the corresponding image after rectification (right) for an example contrail observation at CAB.	28
4.4	Two consecutive annotation time steps five minutes apart at SGP. <i>Left column:</i> original rectified ASI images. <i>Right column:</i> the same images with polygon annotations overlaid in green. Contrails are numbered 1–5 for reference. These contrails are not part of the dataset; the image is included to illustrate the annotation process and the challenges encountered in practice.	29
4.5	Estimation of altitude and age using advected ADS-B waypoints for the first observation of flight UAL42 at 07:13 UTC at CAB on 9 June 2025.	31
4.6	Width estimation for flight AAL2482, fifth observation, SGP, 09/12/2025 15:43 UTC. <i>Left:</i> ASI image with the annotated polygon overlaid in green. <i>Top right:</i> polygon in centerline coordinates; the dashed rectangle is the MRR, grey lines are perpendicular vertex cuts, and the red line marks the maximum-width cut. <i>Bottom right:</i> measured width at each vertex cut position, with the maximum (3.68 km) marked in red and the effective width (2.26 km) shown as a dotted line.	33
4.7	Decomposition of the wind shear vector $\mathbf{S} = (S_u, S_v)$ relative to the contrail axis. The bearing β is measured clockwise from North. The cross-track component $S_{\perp} = S_u \cos \beta - S_v \sin \beta $ displaces ice crystals perpendicular to the contrail and drives lateral spreading; the along-track component S_{\parallel} stretches the contrail lengthwise and does not affect its width.	36

4.8	Geometry of a contrails parallax shift. A contrail at altitude h_{contrail} above site p appears at position p' in the satellite image, displaced horizontally by Δd . The zenith angle θ_z is measured at the contrail between the local vertical and the line of sight s to the satellite.	38
4.9	Example BTM field (C13–C15) for contrail produced by flight AAL2722. The white outline shows the parallax-corrected contrail polygon. The green rectangle is the MRR; the orange dashed and yellow dotted rectangles mark the outer ($\times 5$) and inner ($\times 2$) background boundaries respectively.	39
4.10	Effective pixel footprint geometry. A fixed instrument angular resolution δ maps to $d_{\text{nadir}} \approx H\delta$ at the sub-satellite point (blue). At an off-nadir site viewed at zenith angle θ_z , the same angular aperture projects to $d_{\text{eff}} = d_{\text{nadir}} / \cos \theta_z$ (red).	40
4.11	Ash composites from GOES-East ABI illustrating the two classified sky conditions. <i>Left</i> : clear sky, 1 March 2025 10:00 CST. <i>Right</i> : cirrus, 2 March 2025 08:00 CST.	41
5.1	Monthly count of individually tracked contrails at the SGP site in 2025. Zero counts in May, June, and October reflect gaps in manual event selection rather than zero contrail occurrence.	43
5.2	Hourly distribution of contrail observations at SGP and CAB.	43
5.3	Screenshot from VOCASION for a contrail observation on 22 April 2025 at 13:53 UTC. The annotated contrail polygon is projected onto the Ash image, but the surrounding natural cirrus makes it difficult to determine whether the contrail is genuinely distinguishable from the background. This observation was labeled <i>not visible</i>	45
5.4	Distribution of maximum contrail width for visible (opaque) and not-visible (hatched) satellite observations at SGP and CAB.	46
5.5	Maximum contrail width at the moment of first satellite detection at SGP and CAB sites.	47
5.6	Age at first satellite detection at SGP and CAB.	48
5.7	Contrail formed by flight JBU577 at 09:42 CST, shown at 5-minute timesteps, with the annotated polygon and the maximum width measured denoted at each timestep. In the Ash RGB images, the blue marker denotes the SGP site location, and the timestamp indicates the scan time of the satellite image.	49
5.8	Satellite visibility fraction as a joint function of contrail age and maximum width at SGP. Cell color encodes the fraction of observations in that (age, width) bin that are visible in GOES; numbers give the observation count per cell. Marginal histograms show the distribution of age (top) and width (right) at first visibility.	50
5.9	BTM distribution at SGP for visible, non-visible, and control observations. The visible distribution is shifted to positive values, validating the subjective visibility labels.	51
5.10	BTM contrast versus maximum contrail width at first visibility at SGP	51
5.11	Temporal evolution of the SIA7429 contrail in ASI (top) and BTM (bottom) over 30 minutes. The green dot marks the SGP ASI location for spatial reference.	52
5.12	SKW4822 on 2026-02-07, maximum width 0.73 km, at the SGP site. The blue marker in the middle panel denotes the location of the camera at SGP. Despite being below the nominal 2 km pixel size, the contrail produces a positive brightness temperature difference (BTM) contrast under clear-sky conditions.	53

5.13	BTD contrast distribution at CAB, separated by visibility label and observation stage, with the control group shown for comparison. The separation between visible and not-visible observations is present but weaker than at SGP.	54
5.14	BTD contrast at first visibility of contrails at CAB.	55
5.15	DKH1660 (width 836 m): (a) Cloud Type RGB, contrail visible; (b) BTD at 2 km, contrail absent with negative contrast; (c) HR IR 10.5 μ m at 1 km, contrail resolved.	55
5.16	UAL945 (width 1.305 km) at first visibility: (a) Cloud Type RGB, contrail visible; (b) BTD at 2 km, positive contrast of 0.238 K despite the sub-pixel width ; (c) HR IR 10.5 μ m at 1 km, contrail resolved	56
5.17	DLH453 (width 1.75 km) at first visibility: (a) Cloud Type RGB; (b) BTD image; (c) HR IR 10.5 μ m image	57
5.18	Observed contrail spreading rate versus age at first satellite detection at SGP and CAB. Each point represents one contrail.	58
5.19	Cross-track vertical wind shear versus observed contrail spreading rate at SGP and CAB. Each point represents one contrail.	59
5.20	Mean Brunt–Väisälä frequency versus mean contrail spreading rate per contrail at SGP (left) and CAB (right), colored by ICAO wake turbulence category: dark shading denotes heavy aircraft (MTOW > 136 000 kg); light shading denotes medium aircraft. Dashed lines show linear fits to each group.	60
5.21	Distributions of ERA5 relative humidity over ice (RH_i) across all contrails, separated by whether the contrail became visible in satellite imagery or not, for SGP (left) and CAB (right). The dashed vertical line marks supersaturation with respect to ice ($RH_i = 1.0$).	61
5.22	RH_i distribution at SGP, separated by "clear" and "cirrus" sky condition labels.	62
5.23	Width at first visibility for SGP (left) and age at first visibility (right), split by sky condition.	62

List of Tables

2.1	Overview of satellite sensors used for contrail detection	14
3.1	All-sky imager deployment sites and camera specifications.	18
3.2	GOES-East ABI channels used in this study.	20
3.3	MTG FCI channels used in this study. All channels are from the Full Disk High Spectral Infrared Normal Resolution product (EO:EUM:DAT:0662).	22
4.1	Calibrated parameters for the ASI cameras at SGP and CAB.	26
4.2	Satellite viewing geometry and effective pixel footprint at each ground site. Nadir resolutions are 2 km for the GOES-East Cirrus RGB product and 1 km for the MTG-I1 Cloud Type RGB product.	40
5.1	Summary of the contrail dataset collected at each site.	42
5.2	Variables per polygon observation.	44
5.3	Contrail width statistics at first satellite detection.	47
5.4	Contrail visibility rates by site and sky condition. Percentages are computed within each sky-condition group.	61

Introduction

Aviation's impact on climate change is estimated to represent approximately 3.5% of all anthropogenic climate forcing, as of 2018 [1]. While the emission of carbon dioxide (CO₂) by burning aviation fuel has been extensively researched, it is estimated that the non-CO₂ impacts of aviation comprise about 2/3 of the net radiative forcing [1]. One of these non-CO₂ impacts are condensation trails, or contrails. They form when the hot, moist aircraft exhaust mixes with colder, ambient air, and can persist and spread into contrail cirrus clouds. These human-induced cirrus clouds have a short term but significant climate impact, modifying the Earth's radiation balance by reflecting incoming solar radiation and reducing outgoing radiation. The effect is estimated to be net warming, and as air traffic continues to grow, contrail cirrus is expected to become an increasingly important contributor to human-induced climate change.

Because persistent contrails can only form and grow within ice-supersaturated regions (ISSRs), mitigation strategies have been proposed that aim to reduce contrail climate impact by modifying flight trajectories to avoid these regions, which are generally limited to a few hundred meters in altitude [2]. These contrail-avoidance strategies might be able to reduce the formation of persistent contrails at the cost of small increases in fuel burn. Several contrail modeling approaches have been proposed to predict the regions where contrails will form, and the resulting climate forcing. However, these models can have large uncertainties, and therefore must be extensively validated against observations, as the practical effectiveness of contrail avoidance depends on how accurately contrail formation and persistence can be predicted.

Observations from remote sensing instruments provide a direct and data-driven means of studying contrails in the real world. Geostationary satellite images in particular, offer broad spatial coverage, and are able to monitor contrails over large time periods at intervals of around 10-15 minutes. At the same time, due to coarse spatial resolution and background variability, detecting contrails in these images can be challenging. It has been estimated that the total lifetime of a contrail can be up to three times longer than that observed by a geostationary satellite [3].

Ground cameras can capture contrails as they form, when they are too narrow and optically thin to be seen from geostationary satellites. Using ground cameras, contrails can be attributed to the flights which produced them with high confidence, and then be used to measure contrail age and width, providing information which would be difficult to get from a satellite image. Historically, contrails in ground camera observations have been used to validate contrail prediction models, and

train detection algorithms [4, 5, 6]. Studies have used ground cameras or Low Earth Orbit (LEO) satellites to quantify observability in geostationary satellite images as a function of contrail width, finding that geostationary imagery misses up to 80% of visually identifiable contrails visible from LEO [7], and that visual inspection of geostationary imagery recovers only 48% of those visible from ground-based cameras [6]. However, no study has continuously followed individual contrails from their formation into geostationary satellite imagery, which is necessary to understand at what point in a contrail's lifetime it first becomes detectable, and how that depends on its physical properties and ambient conditions.

The objective of this thesis is to create a dataset of observational data, combining both ground camera and geostationary satellite imagery, to quantify observability in geostationary satellite images as a function of age, width, background and sensor properties. This dataset will allow for detailed information on contrail observability from geostationary satellite images, and can then be used to validate contrail prediction models, and assess the effectiveness of contrail mitigation strategies.

The structure of this work is as follows: Chapter 2 reviews the relevant literature on contrail formation, prediction, and detection, and defines this studies research questions. Chapter 3 describes the instruments and data sources used to construct the dataset. Chapter 4 presents the methodology used to estimate contrail properties, ambient atmospheric conditions, and observability in satellite images. Chapter 5 will present the results of investigating the research questions, and chapter 6 will present the conclusions of the study and recommendations for future work.

Literature Review

This chapter reviews the relevant literature in the field of contrail observation, used to guide the methodology and lay the groundwork for linking ground-based camera images with satellite observations to analyze contrail evolution and observability.

2.1 Schmidt-Appleman criterion

Contrails can form when the warm, moist exhaust from an aircraft mixes with ambient air at cruise altitudes [8]. The conditions necessary for formation are described by the Schmidt-Appleman Criterion (SAC), which can be demonstrated schematically in Figure 2.1. The y-axis is the water vapor partial pressure, the x-axis is temperature, and the two curves represent the saturation pressure with respect to liquid water (blue) and ice (red).

The green line denotes the mixing line, which shows how the exhaust of an aircraft mixes with the surrounding air. The slope is a function of ambient conditions, exhaust conditions and overall engine efficiency, and therefore is specific to the engine and fuel properties of each aircraft. If this mixing line crosses the liquid saturation curve, the air is supersaturated with respect to liquid water, so the water vapor starts to condense on ambient and exhaust particles, which instantly freeze due to the low ambient temperature and become ice crystals. The end point of the mixing line corresponds to the ambient conditions. If the mixing line ends above the ice saturation line, then the surrounding ambient air is super saturated with respect to ice, and the crystals continue to grow, allowing the contrail to persist and gradually spread. Such a region is called an ISSR. Under these conditions, a narrow line-shaped contrail can evolve into contrail cirrus, which may persist for several hours. If the mixing line ends below the ice saturation line, the ice crystals cannot grow and the contrail starts dissipating.

The formation and persistence conditions mean that not all flights produce contrails, and not all contrails that form go on to become long-lived. This distinction is important for predicting where contrails will occur and which of them will grow large enough to affect the Earth's radiative balance.

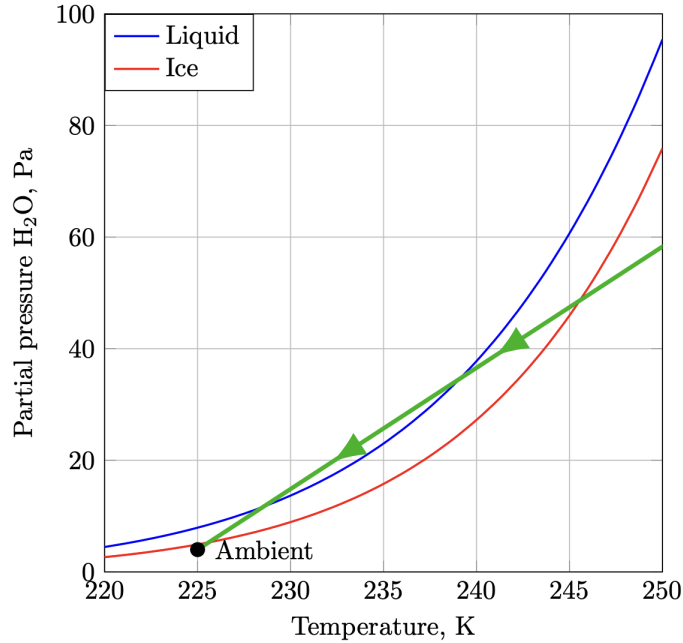


Figure 2.1: Conditions for contrail formation according to the Schmidt-Appleman criterion [9]

2.2 Contrail Cirrus Prediction model

To translate this physical understanding into operational predictions, contrail models have been developed that simulate the formation and evolution of contrails along individual flight trajectories. One of the most widely used is the Contrail Cirrus Prediction (CoCiP), introduced by Schumann [10]. CoCiP predicts where contrails form along a flight path and how they evolve. It combines the contrail formation thresholds given by the SAC with aircraft and engine characteristics and with three-dimensional atmospheric fields from numerical weather prediction models. These meteorological fields provide temperature and humidity at flight altitude, which can be used to determine both whether a contrail can form and whether it will persist in ice-supersaturated regions.

One of the most widely used sources of these meteorological fields in contrail research is the ERA5 reanalysis, produced by the European Center for Medium-Range Weather Forecasts (ECMWF) [11]. ERA5 combines historical observations from satellites, radiosondes, and surface stations with a numerical weather model to produce spatially and temporally consistent estimates of atmospheric variables, available at a horizontal resolution of $0.25^\circ \times 0.25^\circ$ and hourly temporal resolution, and has been used as a common meteorological input for contrail formation and persistence studies [12, 13].

Once a contrail forms, its advection and growth or sublimation are simulated using a Lagrangian Gaussian plume model that represents the contrail as an evolving ice plume. The model uses bulk microphysical parameterizations to describe the ice mass and particle size of the contrail over time. From these simulated contrail properties, CoCiP then estimates the resulting radiative forcing.

Models such as this one can be used for climate impact assessments and contrail avoidance strategies. They allow estimates of contrail climate forcing and can be used to evaluate the potential benefits and costs of rerouting aircraft to avoid ice-supersaturated regions. However, their accuracy depends strongly on the quality of the meteorological input data and on how well the simulated contrail properties (width, optical depth, and lifetime) match reality.

2.3 Satellite-based contrail detection

Satellite imagery has been widely used to gather data on contrail coverage and impact at large scales [[14][15][16]]. However, detecting contrails in satellite images is challenging due to the coarse spatial or temporal resolution of the sensing instruments. While satellites in LEO have higher spatial resolution, they are constrained to two overpasses per day. Satellites in geostationary orbit can provide continuous observations, but their low spatial resolution makes contrails more difficult to detect. An overview of the sensing instruments discussed in this chapter is provided in Table 2.1.

Table 2.1: Overview of satellite sensors used for contrail detection

Sensor	Orbit	Spatial Resolution	Temporal Resolution
AVHRR (NOAA)	Polar / LEO	1.1 km	2 passes per day [17]
MODIS (Terra/Aqua)	Polar / LEO	250m–1 km	2 passes per day [18]
VIIRS (NOAA-20)	Polar / LEO	375–750 m	2 passes per day [19]
CALIOP (CALIPSO)	Polar / LEO	30 m (v) / 5 km (h)	2 passes per day [20]
ABI (GOES-16)	Geostationary	0.5–2 km	10 min [21]
SEVIRI (MSG)	Geostationary	1–3 km	15 min [22]
FCI (MTG-I)	Geostationary	0.5–2 km	10 min [23]

Early satellite-based contrail detection relied on geometric identification of linear features in polar-orbiting imagery, as demonstrated by Mannstein, Meyer, and Wendling [24], who developed the first widely used fully automated Contrail Detection Algorithm (CDA) using NOAA-AVHRR images. The CDA works by recognizing linear-shaped structures in satellite images, using both the brightness temperature difference between the 10.8 μm and 12 μm infrared channels and the absolute brightness temperature in the 12 μm channel to highlight them. Images are screened using line shaped filters in 16 different directions to include all possible orientations, and any structure that fulfills a given a set of geometrical thresholds is labeled as a contrail. Building on this, Vazquez-Navarro, Mannstein, and Mayer [25] applied the CDA to MODIS images to identify linear contrails, and then introduced an Automatic Contrail Tracking Algorithm (ACTA) to follow the same contrails as they evolve in MSG SEVIRI images. In the short timeframe between SEVIRI images, neither the shape nor the position of the detected contrail changes significantly. Gierens and Vázquez-Navarro [3] then used ACTA-tracked contrails in SEVIRI imagery to perform a statistical contrail lifetime analysis, and estimate how much of the contrail life cycle is missed by geostationary sensors due to resolution limits. The contrail lifetime was estimated to be up to three times greater than observed, demonstrating that geostationary satellite observations systematically underestimate both the early and late stages of contrail evolution, due to the contrails being too geometrically narrow or optically thin. This is corroborated by Euchenhofer et al. [7], who compared the same contrails detected in LEO (VIIRS) and GEO (GOES-16 ABI) images and found that the latter fails to recover 80% of the individual contrails and captures less than half of the total contrail length relative to those detected with LEO.

Moving beyond geometry-based CDAs, Meijer et al. [26] trained a deep-learning model to detect contrails on GOES-16 ABI images, allowing for the first automated large-scale analysis of contrail coverage in geostationary satellite images over multiple years. This was expanded by Ng et al. [27], who introduced the OpenContrails benchmark dataset, consisting of over 20 000 human labeled contrail masks, and developed a deep-learning model that incorporated temporal context, which allows the algorithm to better distinguish contrails from natural cirrus or other static linear features by observing their sudden appearance and movement over time. Building on this, Ortiz et al. [28] systematically evaluated multiple neural network architectures on the OpenContrails dataset, finding that inter-annotator disagreement among human labellers introduces an upper bound on achievable detection accuracy, highlighting the inherent subjectivity of contrail annotation as a

limiting factor for detection model performance.

Meijer et al. [29] further extended the detection algorithm of Meijer et al. [26] by collocating detected contrails with CALIOP lidar measurements, resulting in 3D estimates of contrail locations. The resulting dataset was then compared with Automatic Dependent Surveillance-Broadcast (ADS-B) and radiosonde data to assess for plausibility and correctness, and then used to train a new deep-learning algorithm for contrail altitude estimation.

The above approaches demonstrate the ability of geostationary satellite images to provide extremely valuable information on contrail properties, life cycle and coverage, particularly when combined with other instruments and data sources.

2.4 Ground-based camera contrail detection

Several studies have used ground-based cameras to obtain higher-resolution, continuous observations of contrails, which can be used to study the early lifetime and width of contrails with higher fidelity, providing us with valuable information on contrail evolution.

Mannstein, Brömser, and Bugliaro [6] used all-sky cameras to detect and determine the width of specific contrails, and used these measurements to quantify the detection efficiency of the CDA on NOAA-AVHRR and MSG-SEVIRI images as a function of width. Their method consists of collecting all contrails with a width of 1-5 km from the all sky camera, and comparing these with the output of contrails detected using the CDA as well as contrails detected by visual inspection of the satellite images. The findings show that for MSG SEVIRI images, visual inspection yields 48% of the contrails of 1-5 km width, and the CDA 19%.

Another advantage of ground-based cameras is in using their observations to validate formation predictions made by the SAC, or simulations of contrails created by CoCiP. Schumann et al. [4] used 4 fixed wide-angle cameras placed several kilometers apart to determine the three-dimensional position of specific contrails. These contrails were then used to identify the aircraft which caused them through comparison with ADS-B data, and also used to compare with contrails simulated by CoCiP [10]. In a similar vein, Low et al. [13] used ground-based camera observations to compare the simulated predictions of contrail formation by SAC and CoCiP from ADS-B aircraft trajectories to the observed contrails formed above the camera by the same trajectories. The observed lifetime and width was also compared to the simulated lifetime and width of CoCiP. Discrepancies between simulated and observed contrail characteristics in these studies arise from multiple sources, including uncertainties in the ERA5 meteorological fields used as model input, unresolved sub-grid-scale variability at flight altitudes, simplifying assumptions within the contrail models themselves, and uncertainties in the representation of aircraft engine properties.

Ground-based camera observations have also been used to create datasets for training contrail detection algorithms. Gourgue, Boucher, and Barthès [5] introduced an annotated dataset of all sky images, with labels for various cloud conditions and contrail ages, ranging from young linear contrails to older, diffuse ones. Jarry et al. [30] presented the larger Ground Visible Camera Contrail Sequences (GVCCS) dataset, consisting of over 24 000 annotated frames with instance-level annotations, and presented a deep learning method in which each contrail is individually labeled, tracked over time and linked to specific flights through ADS-B data, allowing a detailed analysis of its early life cycle. Together these datasets provide a foundation for training and benchmarking ground-based contrail detection and tracking algorithms.

2.5 Flight to contrail attribution

Several recent approaches aim to attribute specific contrails observed using remote sensing instruments to the flights that produced them. This is important for validation of contrail models, where knowing the aircraft and engine type, as well as altitude flown and ambient conditions nearby, can enable a better comparison between the model prediction and observations. In most cases, these methods follow a similar sequence: contrails are first detected in satellite images, then candidate flight trajectories are advected forward in time using wind data, and finally matched to candidate flight trajectories based on spatial and temporal alignment. Modern implementations of this pipeline rely on computer-vision-based contrail detection combined with ADS-B trajectories and meteorology-driven advection models ([31], [32]). These methods are limited by several factors, such as high air traffic densities, errors in wind data and delay between contrail formation and observability in satellite images. Within this general framework, different strategies have been proposed.

Chevallier et al. [31] presented a framework for contrail tracking and flight attribution in geostationary satellite imagery, in which detected contrails are linked to aircraft trajectories through a global optimisation approach. Contrail detection was performed using a deep-learning-based algorithm applied to satellite imagery, while tracking and attribution are formulated as a joint problem that incorporates multiple data sources, including ADS-B trajectories, ERA5 wind fields, and CoCiP-simulated contrails. As no open labeled dataset existed at the time, validation relied on qualitative case studies of isolated contrails. Geraedts et al. [32] implemented an Automated Detection and Matching (ADM) system, a pipeline that automatically detects contrails in satellite imagery and matches them to candidate flight trajectories, applying this approach at large scale to hundreds of thousands of flights and compared observations directly with SAC and CoCiP predictions. Validation against a human-labeled subset of 1,000 flight segments showed reasonable performance in low-traffic conditions, though dense airspace remained a consistent challenge.

In order to benchmark and improve these contrail-to-flight attribution algorithms in a consistent way, Sarna et al. [33] introduced a synthetic, flight-attributed contrail dataset, together with a new attribution algorithm. This study provided a structured comparison of the different approaches discussed above, and defined quantitative metrics for evaluating attribution performance. Using this benchmark, they showed how contrail properties such as altitude, density, age, and length influence how reliably contrails can be matched to flights. This study shows the need for a comprehensive evaluation of detection and attribution algorithms against real-world ground-truth data.

Gryspeerd et al. [34] matched air traffic data to satellite observations to examine how contrail formation and lifetime depend on aircraft type, cruise altitude, and atmospheric conditions, finding a strong link between more efficient aircraft flying at higher altitudes and higher SAC temperature thresholds and longer living contrails.

Driver, Stettler, and Gryspeerd [35] developed a framework to quantify the theoretical detectability of contrails in geostationary infrared satellite imagery. They simulated contrail segments as straight-line clouds with a Gaussian ice water path profile, using radiative transfer simulations to produce brightness temperature fields at satellite resolution, and systematically varied the contrail width, depth, altitude, ice water path and effective radius to quantify which properties affected detectability the most. The observability threshold derived from these simulations was then applied to a population of contrails modeled by CoCiP, providing a realistic distribution of contrail properties against which to assess observability. It was found that the properties with the strongest observability responses were the contrail width, effective radius, which describes the size of constituent ice crystals, and ice water path, which describes the mass of ice in a contrail per unit horizontal area. It was also found that background properties such as a layer of natural cirrus strongly limits

observability. For an infrared image with a resolution of 0.5 km under "best-case" conditions, about a third of the contrail population was found to be unobservable. However, their study used synthetic contrails simulated on idealized scenes over a uniform ocean background, which does not capture the variability of real atmospheric and surface conditions. This highlights a remaining gap in quantifying how contrail properties, background conditions, and instrument characteristics together control observability in real satellite imagery.

2.6 Research gap

Despite the advances described above, satellite-based and ground-based contrail observations have largely remained separate: satellite-based methods begin only once contrails are visible from space, and ground-based observations are rarely followed into the satellite images. As a result, there is still no framework that provides a continuous, flight-attributed view of contrail evolution from initial formation to its later satellite-observed stages. There is a need for a comprehensive labeled contrail dataset, following attributed contrails from formation into spreading and dissipation. Furthermore, it is still unclear how the combination of contrail width, lifetime, the surrounding cloud and surface background, and the properties of the satellite sensor, together control whether a contrail appears in satellite imagery and how long after its formation it can first be detected.

Accordingly, this thesis will address these concerns with the following research questions:

- 1. Which combination of contrail properties, background conditions, and instrument characteristics determines whether a particular contrail is observable in a given satellite image?
- 2. How long does it take for a contrail to become observable in satellite imagery, and how does this delay depend on atmospheric and background conditions?

Data Description

Addressing the research questions requires observations that bridge the gap between contrail formation at the ground and first detection from space. This chapter describes the instruments and data sources used for that purpose, combining ground-based all-sky imagery, geostationary satellite observations, meteorological reanalysis, and flight trajectory data.

Ground-based observations were collected at two sites: the ARM Southern Great Plains (SGP) facility in Lamont, Oklahoma, USA [36], and the KNMI Cabauw (CAB) Experimental Site for Atmospheric Research (CESAR) in the Netherlands [37]. The entire data collection was facilitated by VOCASION, an application developed to view and label contrails across both ground-based and satellite imagery on both sites. The instruments and data sources are described in Sections 3.1–3.4, and VOCASION is presented in section 3.5.

3.1 All-Sky Imager

At each site, an upward-facing fisheye All Sky Imager (ASI) captures images of the overhead sky at one-minute intervals. The two deployment sites and their camera specifications are summarized in Table 3.1.

Table 3.1: All-sky imager deployment sites and camera specifications.

	SGP	CAB
Full name	ARM Southern Great Plains	CESAR Observatory
Location	Lamont, Oklahoma, USA	Cabauw, Netherlands
Altitude (above sea level)	314 m	−0.7 m
Latitude	36.607° N	51.97° N
Longitude	97.488° W	4.93° E
Instrument	All-Sky Imager (ASI)	All-Sky Imager (ASI)
Native resolution	1536 × 1536 px	1920 × 1920 px

These ASI images form the basis of the entire dataset. For this dataset to be useful, it requires a large number of persistent contrails that are likely to become visible in geostationary satellite imagery.

In order to efficiently collect as many contrail sightings as possible within a limited amount of time, the following method was used to identify the best dates to analyze images for.

3.1.1 Contrail Activity Time Series and Case Selection

From the Google Contrails Explorer [38], which applies the deep-learning detection algorithm of Ng et al. [27], a time series of contrail detection counts above the measurement site was obtained, aggregated into 10-minute bins. To identify dates with sustained contrail activity, the time series was processed using connected component labeling: consecutive 10-minute bins in which at least one contrail was detected above the site were grouped into events. An event is defined as a sequence of such bins with no gap. A single bin with no detections breaks the sequence and starts a new event, as visualized in Figure 3.1.




Figure 3.1: Illustration of the event-definition rule. Each rectangle represents one 10-minute bin. Blue bins contain at least one contrail detection; gray bins contain none. A sequence of consecutive blue bins with no intervening gap forms a single event (Event 1, Event 2, Event 3). A single grey bin is sufficient to terminate the current event and start a new one.

Events were then aggregated by date, yielding for each day the number of distinct events and the total contrail activity duration. A day in which 10 "events" of contrail detections above the site occurred spanning all daylight hours, for example, might be a very good day to analyze. By visually verifying each day and time of interest on the satellite viewer in VOCASION (see section 3.5), the specific dates and times could be selected for further analysis.

In the end, 97 hours of images, over 29 days across 2025 and the start of 2026 in SGP were analyzed. In CAB, 58 hours of images across 11 days in 2025. In total, 155 hours of ASI images were manually analyzed to collect contrail observations for this dataset.

3.2 Satellite Imagery

3.2.1 GOES-East ABI

For contrails above the SGP site, the GOES-East satellite, operated by NOAA, is used. GOES-East is positioned at 75° W longitude at an altitude of 35 786 km above the equator. During the earlier portion of the dataset this slot was occupied by GOES-16, and by GOES-19 from April 2025 onwards. Both carry the ABI, a 16-channel passive imaging radiometer [39]. Unlike a typical ground-based camera, the ABI does not capture an instantaneous snapshot of what it sees in its field-of-view (FOV). Instead, it scans the Earth by sweeping a detector array in a series of east-west scan lines, called swaths, stepping progressively from north to south. The FOV of the ABI can be seen in Figure 3.2a. A full disk scan takes approximately 9 minutes and 31 seconds, meaning that pixels at different latitudes within a single 10-minute image are acquired at different times, as illustrated in Figure 3.2b. The swaths for a full disk scan are shown in pink, with each new swath being scanned

approximately every 30 seconds.

The ABI records the SGP site in Oklahoma approximately 2–3 minutes into the full disk scan. This is relevant because contrails advect with upper troposphere wind speeds: within a single 10-minute image interval, a contrail can form, drift significantly, and evolve in width and optical depth. The temporal offset between the ASI observation and the satellite observation therefore introduces a positional uncertainty that must be accounted for when tracking the contrail from the ASI into the satellite image.

In addition, because GOES-East observes the Earth from a fixed position at 75° W, 0° N, any location away from the sub-satellite point is viewed at an oblique angle. For features at altitude such as contrails, this introduces a parallax displacement between the apparent and true geographic position, the correction of which is discussed in Section 4.3.1. Four ABI channels are used in this study, listed in Table 3.2.

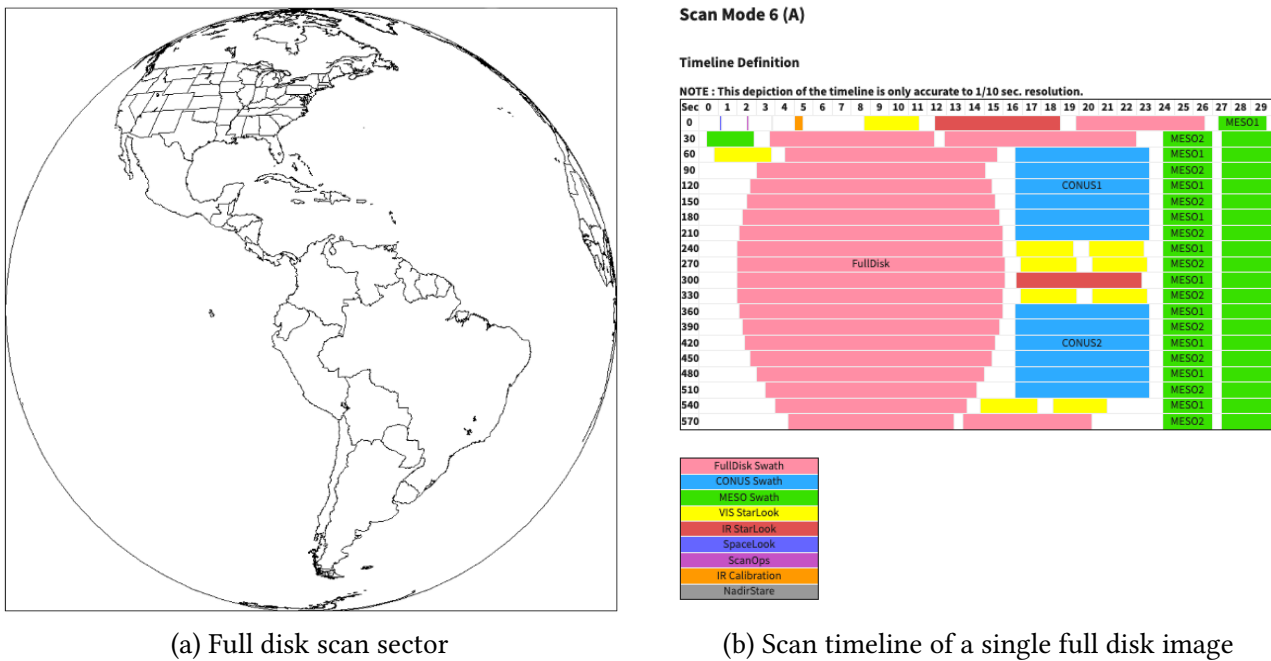


Figure 3.2: GOES-East ABI Mode 6 scanning: spatial coverage and timeline [40]

Table 3.2: GOES-East ABI channels used in this study.

Channel	Central wavelength	Nadir resolution	Use
11	8.4 μm	2 km	Ash RGB
13	10.3 μm	2 km	BTD
14	11.2 μm	2 km	Ash RGB
15	12.3 μm	2 km	Ash RGB, BTD

In this study, two satellite products are used, serving different purposes. The Ash RGB is used by VOCASION (see section 3.5) for manual, visual labeling, serving as the primary method for determining contrail visibility in the dataset. To provide an objective verification of these manual labels, a pixel-level BTD analysis is performed on each annotated contrail polygon. Both products are derived from ABI infrared channels listed above.

Manual Detection: Ash RGB

For manual, visual detection of contrails, the Ash RGB false-color composite is used. This product leverages three thermal infrared bands: C11, C14, and C15, with red representing the $12.3\ \mu\text{m}$ band, green the difference between $12.3\ \mu\text{m}$ and $11.2\ \mu\text{m}$, and blue the difference between $11.2\ \mu\text{m}$ and $8.4\ \mu\text{m}$ [41]. This specific combination renders contrails as dark blue linear features, making them easier to identify. Historically, Ash RGB images have been used to create manually labeled contrail datasets, such as in Ng et al. [27] and Sarna et al. [33]. An example of this RGB can be seen in Figure 3.4, as the image on the right.

Pixel-Level Analysis: Brightness Temperature Difference

For pixel-level verification, the Level-1b full-disk radiance product (ABI-L1b-RadF, Mode 6) is processed to generate the BTD images. Brightness temperature is the equivalent blackbody temperature that would produce the observed thermal radiance at a given wavelength. Subtracting the brightness temperature of the band centered at $12\ \mu\text{m}$ from the one centered at $10\text{--}11\ \mu\text{m}$ highlights thin cirrus, rendering contrails in BTD images as bright features characterized by large, positive values. This is called an infrared split-window technique, and has been historically used for analyzing contrails in images [42, 43, 15, 34]. The BTD image also renders any pixels without cloudiness, such as land or water, as values near zero. This means that the contrast of the contrail from its surroundings will provide an objective verification that the annotated contrail polygon possesses the specific radiative signature of a thin ice cloud.

3.2.2 MTG FCI

For the CAB site, imagery from the Meteosat Third Generation Flexible Combined Imager (FCI) is used. MTG-I1 is positioned at 0 E longitude at an altitude of approximately 35 786 km. The Level-1C netCDF product is used [44].

Like the GOES-East ABI, the FCI does not capture an instantaneous snapshot of the full disk. Instead, it scans the Earth in a series of east-west and west-east alternating swaths, stepping progressively from south to north. The full disk is completed in approximately 9 minutes and 30 seconds. The CAB site is located in the northern part of the full disk, and is scanned at approximately the 8-minute mark.

For manual detection of contrails, the Cloud Type RGB is used, which can separate different types of clouds successfully. The $1.38\ \mu\text{m}$ red channel allows for detection of very thin cirrus clouds over other clouds and land [45], making it excellent for contrail observations. For pixel level analysis of each contrail, the brightness temperature difference between the $10.5\ \mu\text{m}$ and $12.3\ \mu\text{m}$ channels will be used. The overview of all channels used in this study are shown in Table 3.3.

3.3 ERA5 Reanalysis

Meteorological fields are obtained from the ERA5 reanalysis produced by the European Centre for Medium-Range Weather Forecasts (ECMWF). The model-level product [46] is retrieved at $0.25^\circ \times 0.25^\circ$ horizontal resolution and 1 h temporal resolution. Temperature, specific humidity and horizontal wind components (u , v) are requested on model levels 59–85, covering approximately 8.5–14 km altitude. The methods used to derive ambient conditions around each contrail observation

Table 3.3: MTG FCI channels used in this study. All channels are from the Full Disk High Spectral Infrared Normal Resolution product (EO:EUM:DAT:0662).

Channel	Central wavelength	Nadir resolution	Use
03	0.6 μm	1 km	Cloud Type RGB (Green)
06	1.38 μm	1 km	Cloud Type RGB (Red)
07	1.61 μm	1 km	Cloud Type RGB (Blue)
14	10.5 μm	2 km	BTD
15	12.3 μm	2 km	BTD

from these fields are described in section 4.2.

3.4 ADS-B Flight Data

ADS-B flight data is sourced from the contrails.org ADS-B API, and filtered to only waypoints passing over a $\pm 1.5^\circ$ bounding box centered on each camera site. On its own, a raw flight track shows where an aircraft was at the moment it flew overhead, but by the time a contrail is observed in the ground camera, the wind may have carried it away from that original position. To correctly identify which flight produced a given contrail, the flight tracks must be advected forward in time to match the image timestamp.

3.4.1 Flight Track Advection

For each ASI image time t , all ADS-B waypoints recorded within the preceding hour are advected forward to t using the AFCA library [9], a high-performance C++ implementation of Lagrangian advection developed for contrail analysis. This advection is performed within VOCASION (see section 3.5) and the resulting trajectories are made available for download as `.csv` files, one per image timestamp. For each waypoint, the file records the original flight position and timestamp, the advected position at t as longitude and latitude, and the corresponding projected pixel coordinates (u, v) in the rectified all-sky image. As shown in Figure 3.3, the advected tracks are overlaid in green on both the ASI and satellite images, allowing the annotator to visually identify which flight track aligns with each visible contrail. These precomputed files are also used in the processing pipeline for altitude and age estimation, as described in Sections 4.1.6.

3.5 VOCASION Description

VOCASION is the application used to perform the entire contrail database collection. It is a `Next.js` application that enables the user to view and annotate contrails in ground-based ASI and satellite images simultaneously. In this section, the principal features of VOCASION will be presented. An image of the main UI is shown in Figure 3.3, with each annotated red number describing a feature as explained in the list below.

1. Menu bar to explore the different functionalities. Tracking and annotating contrails is done in the current view, called the *Map* view. The *Calendar* shows an overview of all dates for which ASI images have been downloaded, *Download* allows the user to download the image

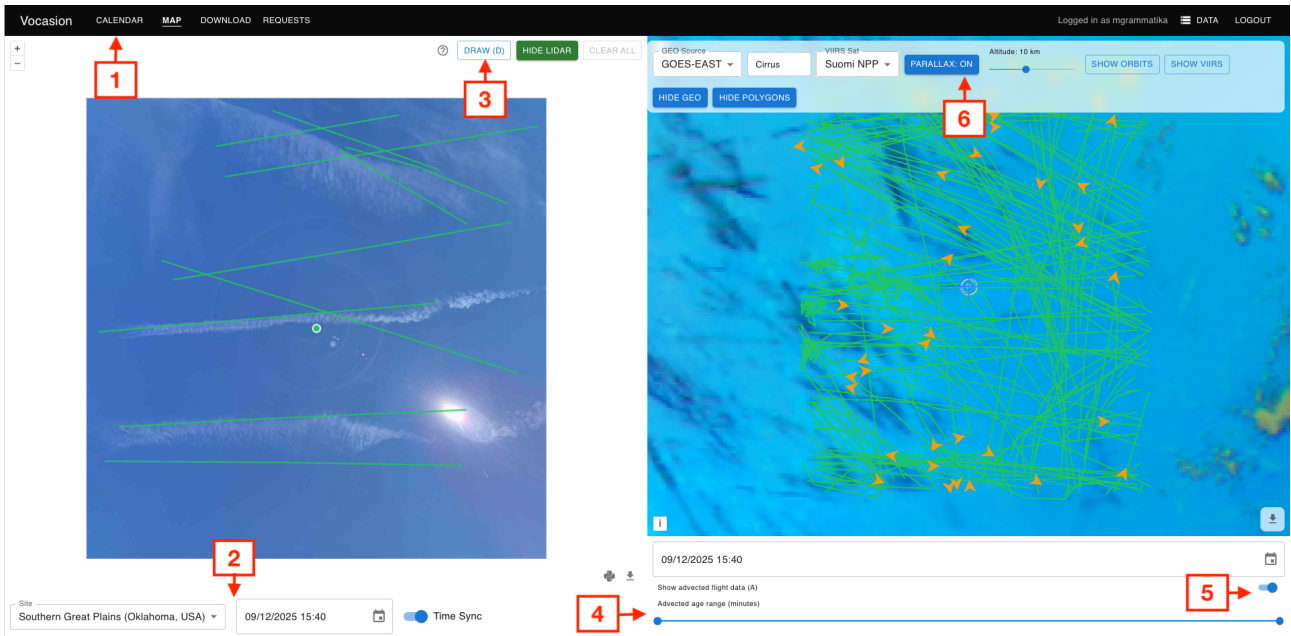


Figure 3.3: VOCASION UI

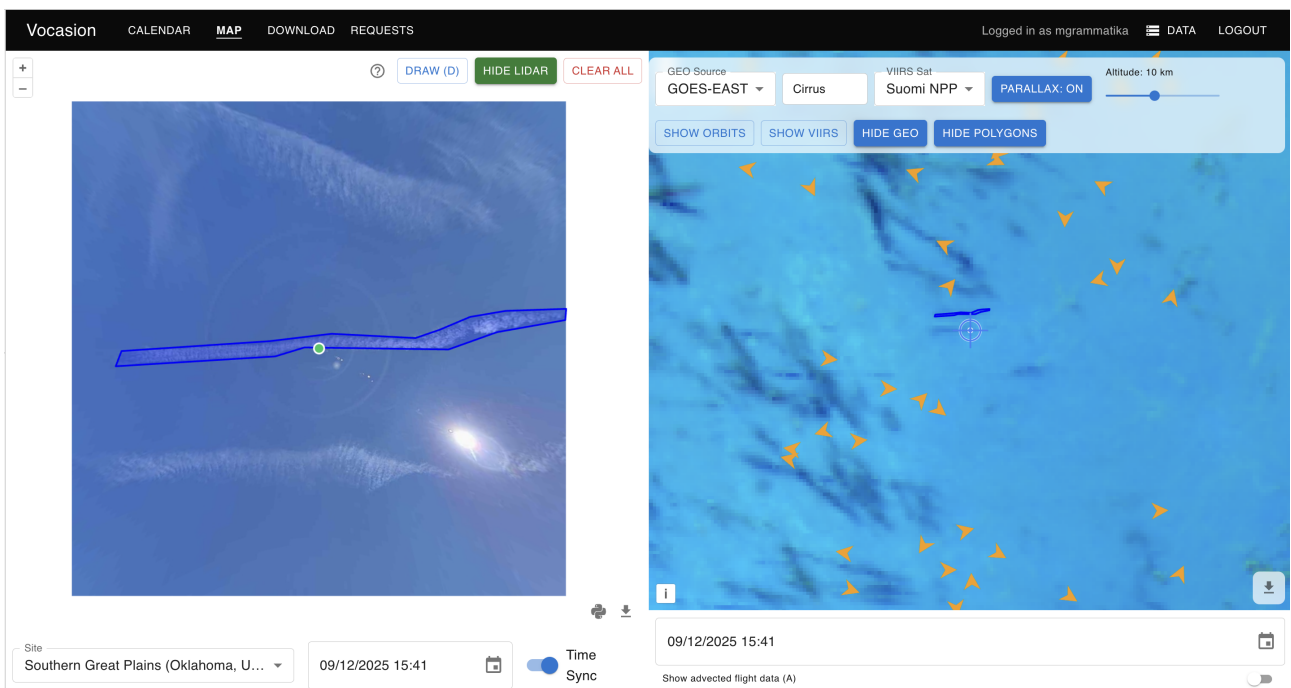


Figure 3.4: VOCASION polygon example

and flight data at a specific datetime. The *Requests* feature can be used to request new ASI images or advected flight data.

2. Drop-down field to select the site to look at. Date selection field to step through the images. Time sync button syncs the ground camera timestamp with the satellite image immediately preceding. For example, an ASI image at 15:49 will be matched to the satellite scan beginning at 15:40.
3. "Draw" button used to start drawing a polygon over a contrail observation. An example of a polygon can be seen in Figure 3.4. Note how VOCASION automatically projects the same polygon into the satellite image, for matching.

4. Slider to adjust length of advection of each waypoint. This can be used to limit the number of flights visible overhead to only the most recent ones.
5. Toggle to turn the advected flight tracks on and off.
6. Toggle to turn parallax shift on and off. The slider shifts from 7-14 km, and is used to move each contrail polygon drawn in the ASI to the position where it shows up in the satellite. Parallax shift is explained in further detail in subsection 4.3.1.

Methodology

This chapter describes the methods used to process the dataset of attributed contrails tracked from formation into visibility in satellite images. For each contrail polygon annotated in the ASI imagery, the width and age at the time of observation are estimated, as described in section 4.1. The ambient meteorological conditions at contrail altitude are then interpolated from ERA5 reanalysis fields, as described in section 4.2. Finally, each polygon is projected into the corresponding satellite image and a brightness temperature contrast is extracted to provide an objective measure of satellite visibility, as described in section 4.3.

4.1 Contrail Width and Age Estimation

Estimating the physical width and age of a contrail from a ground-based fisheye camera involves several steps: modeling the lens geometry, rectifying the image, and matching the annotated polygon to an ADS-B waypoint to obtain the contrail altitude and age. The altitude is then used to scale the polygon to physical distances, from which the width is measured along the contrail axis. Each step is described below.

4.1.1 All-Sky Imager Setup

The ASIs used in this study are deployed at two sites: the Southern Great Plains (SGP) ARM Climate Research Facility in Lamont, Oklahoma, USA (36.607° N, 97.488° W) [36], and the Cabauw Experimental Site for Atmospheric Research (CESAR) in the Netherlands (51.968° N, 4.928° E)[37]. Both cameras are equipped with fisheye lenses that provide a 180° view of the sky.

Each ASI camera is characterized by a set of calibration parameters, listed in Table 4.1, which describe the lens geometry, where $f_x = f_y = f$ is the focal length in pixel units and (x_0, y_0) is the principal point. The parameter ψ denotes the rotation of the camera about its vertical axis, measured counter-clockwise from the local East direction. These calibrated parameters are used throughout the projection pipeline described in the sections below.

4.1.3 Image Rectification

Because the fisheye projection is non-linear, contrails, which are characteristically linear features, appear as curves in the raw image. For accurate detection and analysis, contrails must be represented as straight lines. The fisheye image is therefore transformed into a rectilinear image through a process called rectification.




Figure 4.2: Pinhole camera model. A real-world point $P(X, Y, Z)$ is projected through the camera onto the image plane at pixel $p'(u, v)$, at radial distance r from the principal point (u_0, v_0) and viewing angle θ_p from the optical axis.

Under the pinhole camera model, straight lines in the world remain straight lines in the image. Fisheye lenses do not follow this model, so contrails appear as curves. Since linearity is a defining property of contrails, this must be corrected before any detection or analysis can be performed. Rectification maps the fisheye image to a rectilinear one, restoring straight-line geometry. A rectilinear image follows the pinhole camera model, illustrated in Figure 4.2, in which a real-world point $P(X, Y, Z)$ is projected onto the image plane at pixel $p'(u, v)$. The radial distance from the principal point (u_0, v_0) to the projected pixel is:

$$r = \sqrt{(u - u_0)^2 + (v - v_0)^2}, \quad (4.3)$$

and the corresponding viewing angle from the optical axis is:

$$\theta_p = \arctan\left(\frac{r}{f_R}\right). \quad (4.4)$$

Rectification works by inverse mapping: for each output pixel, θ_p is computed from Equation 4.4, and the distortion model then gives the corresponding source pixel in the fisheye image. The rectified image is defined by a maximum viewing angle of $\theta_{p,\max} = 60^\circ$ from zenith, mapped onto a 901×901 px output. The decision to use a cut-off angle of 60° is in line with the recently published dataset of contrails labeled in all-sky imagery by Gourgue, Boucher, and Barthès [5]. The rectilinear focal length f_R is chosen so that a ray at exactly $\theta_{p,\max}$ lands at the edge of the image:

$$f_R = \frac{H_R/2}{\tan \theta_{p,\max}} = \frac{450.5}{\tan 60^\circ} \approx 260.1 \text{ px}, \quad (4.5)$$

with the principal point at the image centre $(u_0, v_0) = (450.5, 450.5)$. An example of a rectified image is shown in Figure 4.3.




Figure 4.3: Raw fisheye image (left) and the corresponding image after rectification (right) for an example contrail observation at CAB.

4.1.4 Contrail Polygon Annotation

Contrail boundaries are delineated manually within VOCASION by clicking vertices to form a closed polygon around the visible contrail in the rectified 901×901 px image. Annotation begins at the moment of first visibility and is repeated at approximately five-minute intervals, producing a time series of polygons that captures the evolution of the contrail boundary over its observable lifetime. Annotation ceases when the contrail either advects out of the camera FOV or becomes too diffuse to delineate meaningfully.

Figure 4.4 shows two consecutive annotation time steps, five minutes apart, with the original and annotated images, illustrating the range of challenges encountered in practice. For reference, the visible contrails are numbered 1–5 in the annotated frames. In the top left frame, several contrails are visible simultaneously at different stages of development. Contrail [1] is diffuse and partially transparent but its boundary is still just distinguishable from the background, and a polygon is drawn around it. Contrail [2] has developed an irregular shape, which the annotation attempts to follow closely. Natural cirrus is present in parts of the image, making some edges harder to place precisely.

Five minutes later, shown in the bottom two frames, the eastward winds have advected all contrails further across the frame. Contrail [5] has left the FOV entirely and receives no further annotation. Contrail [1] has become too diffuse for its boundary to be delineated with confidence, and is therefore also not annotated. Contrail [3] has gone from having just been formed to now 5 minutes later having spread into an irregular shape. The remaining contrails have all spread and changed shape; each is re-annotated with a new polygon reflecting its current extent. In all cases, the goal of the annotation is to capture the widest points of the contrail within the boundary and to use as many vertices as necessary to follow the visible shape. This is inherently subjective, particularly for irregular or diffuse boundaries, and the consequences for width estimation are discussed in

subsection 4.1.8. Note that the contrails shown are not part of the dataset assembled in this study, and this example is included purely to illustrate the annotation process.

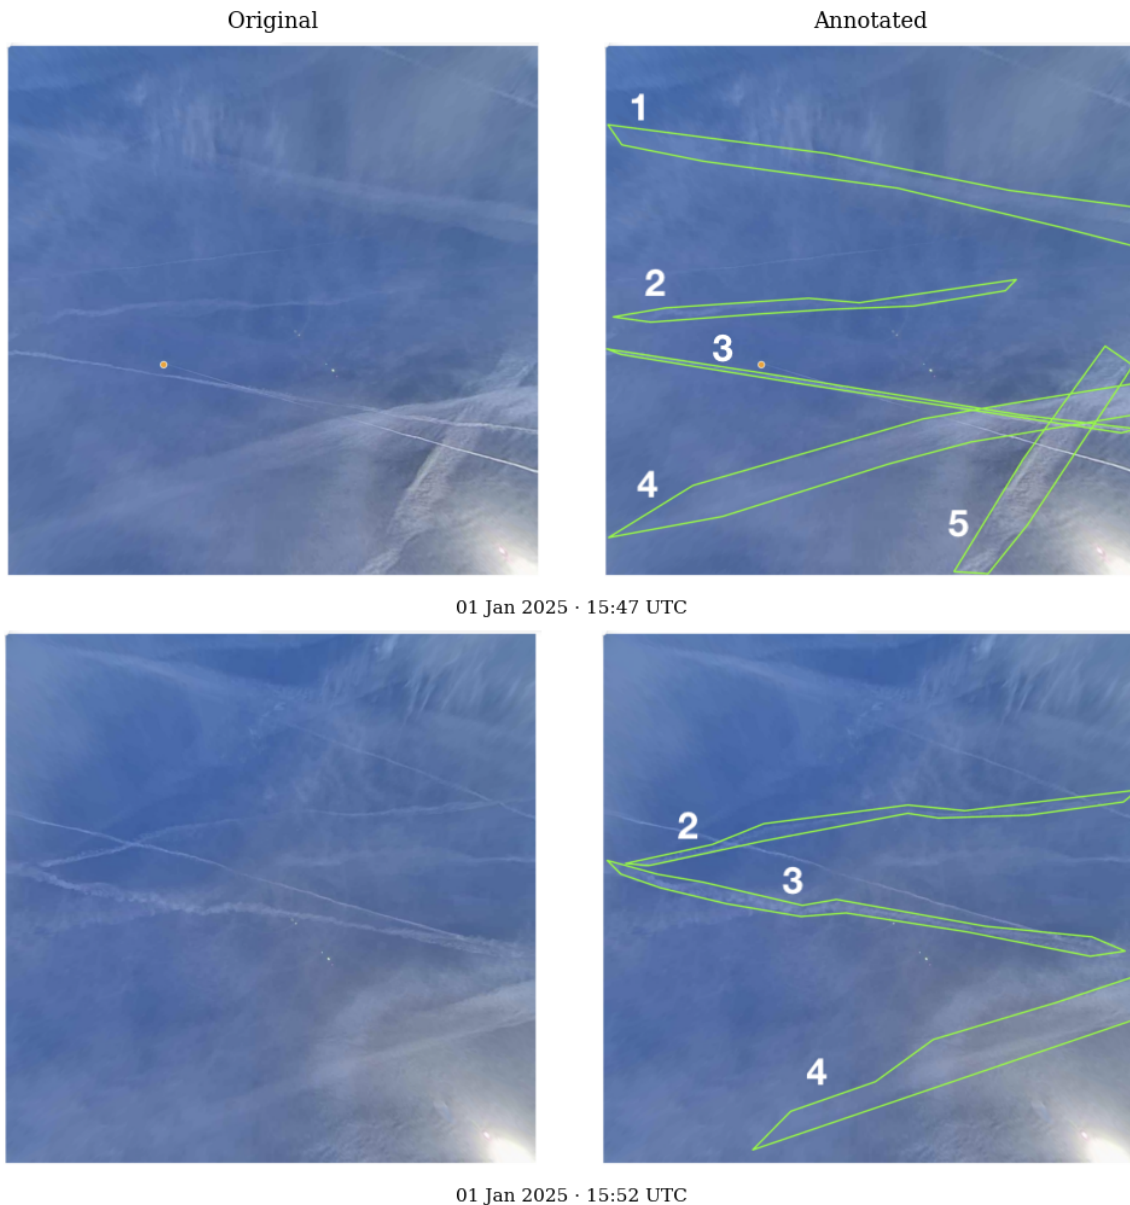


Figure 4.4: Two consecutive annotation time steps five minutes apart at SGP. *Left column*: original rectified ASI images. *Right column*: the same images with polygon annotations overlaid in green. Contrails are numbered 1–5 for reference. These contrails are not part of the dataset; the image is included to illustrate the annotation process and the challenges encountered in practice.

The key point is to show that boundary placement is inherently subjective. Different annotators, or even the same annotator on different occasions, may draw slightly different polygons for the same contrail, particularly where the edge is diffuse or the background is complex. This subjectivity propagates directly into the width measurements described in subsection 4.1.7, and is quantified in subsection 4.1.9.

4.1.5 Projection to Physical Coordinates

With the contrail polygon defined in the rectified image, the next step is to convert those pixel coordinates into physical distances. This requires knowing the altitude of the contrail, as the same

angular size in the image corresponds to a larger physical width at higher altitude.

Each rectified pixel (u, v) corresponds to a unique direction in space, as illustrated in Figure 4.2. Projecting onto a plane at unit height $Z = 1$ above the camera gives the unit-height coordinates:

$$X = \frac{u - u_0}{f_R}, \quad Y = \frac{v - v_0}{f_R}, \quad (4.6)$$

giving a dimensionless polygon $\{(X_i, Y_i)\}_i$ whose shape is independent of altitude and can be scaled to any altitude in a single step.

Altitude Scaling

To recover physical distances, the unit-height polygon is scaled by the contrail's height above the camera:

$$h_{\text{rel}} = h_{\text{contrail}} - h_{\text{cam}}, \quad (4.7)$$

where h_{contrail} is the contrail altitude in meters above mean sea level (MSL) and h_{cam} is the camera altitude MSL (314 m for SGP; -0.7 m for CAB). Each vertex (X, Y) of the unit-height polygon maps to physical coordinates in meters via

$$x_m = X \cdot h_{\text{rel}}, \quad y_m = Y \cdot h_{\text{rel}} \quad (4.8)$$

4.1.6 Altitude and Age Estimation

As explained in Section 3.4.1, each ADS-B waypoint of a flight in an ASI image above the site is advected forward in time to the observation timestamp, so that each waypoint reflects where that parcel of air would be at the moment of the image. Matching the observed contrail to the waypoint that produced it simultaneously yields the contrail altitude and age.

For every contrail in the dataset, the centroid of the annotated polygon at the first observed time step is computed as the area-weighted geometric centroid of its vertices using the Shapely library [49]. The centroid is then matched to the nearest advected waypoint in pixel coordinates:

$$d = \sqrt{(u_w - C_u)^2 + (v_w - C_v)^2}, \quad (4.9)$$

with (u_w, v_w) the projected position of waypoint w . The geometric altitude of that waypoint, representing the altitude at which the aircraft was flying when it produced that contrail segment, is taken as h_{contrail} and applied uniformly to all subsequent time steps. The contrail age at first observation is then:

$$t_0 = T_{\text{obs}} - T_{\text{dep}}, \quad (4.10)$$

where T_{obs} is the observation time and T_{dep} is the original timestamp of the matched waypoint. For all subsequent frames, the age increases by the elapsed observation interval:

$$t_n = t_0 + (T_n - T_0), \quad (4.11)$$

where T_n is the ASI timestamp of frame n and T_0 that of the first frame. Polygons are drawn at approximately five-minute intervals, so age increases in five-minute steps across the observation window.

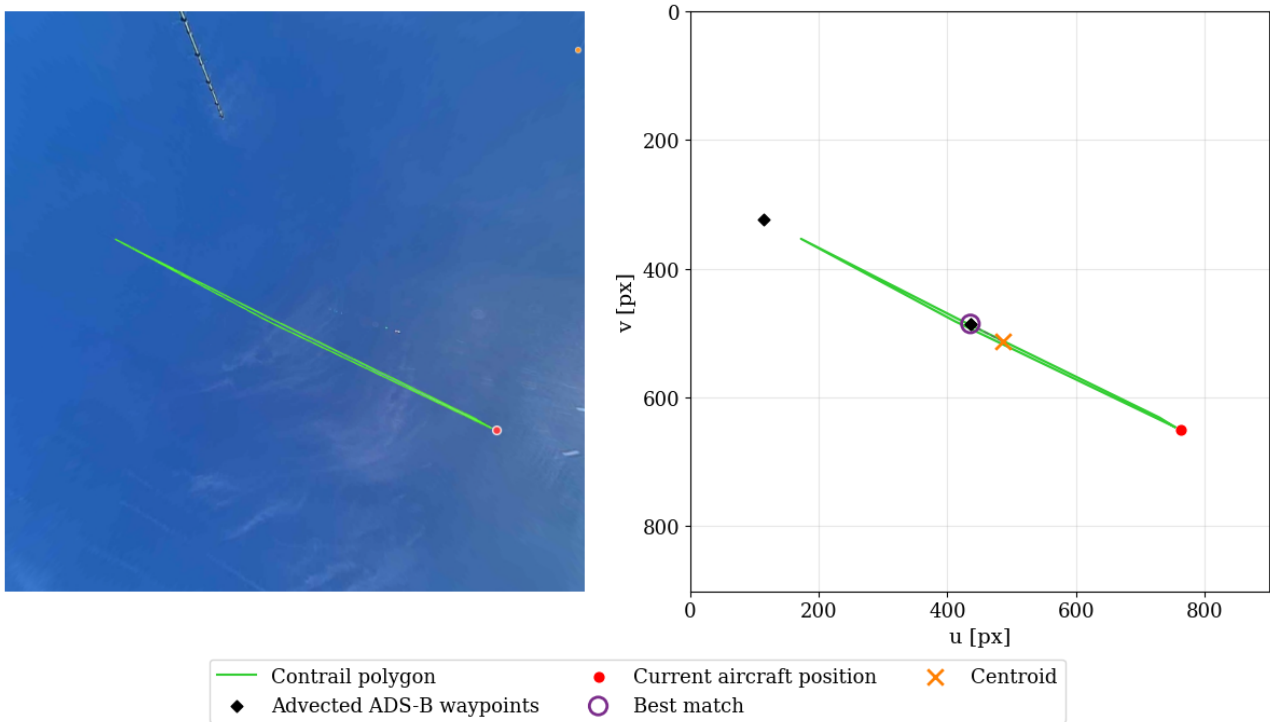


Figure 4.5: Estimation of altitude and age using advected ADS-B waypoints for the first observation of flight UAL42 at 07:13 UTC at CAB on 9 June 2025.

This methodology is illustrated in Figure 4.5 for the first observation of flight UAL42. On the left, the ASI image shows the polygon as annotated in VOCASION in lime green, and the current aircraft position as the red circle. On the right, the same polygon is shown in pixel coordinate space. The grey diamonds show the advected ADS-B waypoints of the two previous minutes of flight history. The centroid of the annotated polygon (orange cross) is matched to the nearest waypoint (purple circle), which lies one minute earlier than the current aircraft position and has therefore been advected one minute forward in time, giving a contrail age of 1 minute. The geometric altitude of that waypoint is assigned as $h_{\text{contrail}} = 10.6$ km.

Limitations

Sources of uncertainty affect the altitude and age estimates produced by this method. Once the match is established, a single geometric altitude taken from the first observation of the contrail in the ASI, is applied uniformly to all subsequent time steps. This assumes that the annotated polygon tracks the same parcel of air across all frames, which may well hold for short observation windows or slow-moving contrails. In practice, however, it is not guaranteed: as the contrail advects, parcels visible at first observation may drift out of the camera field of view while new parcels, formed at a different time or altitude along the flight track, enter the frame and become part of the annotation. In such cases, the age and altitude assigned from the first-frame match might become increasingly unrepresentative of what is actually being observed.

4.1.7 Width Estimation

Contrail width is not uniform along its length: a contrail can taper, bulge, or develop irregular boundaries as it spreads. The following subsections describe how a longitudinal axis is estimated from the polygon geometry and how cross-sectional width is measured along it, yielding the two

metrics used in this study. The procedure is illustrated in Figure 4.6 for flight AAL2482 at SGP, which by its fifth observation has developed a distinctly irregular shape.

Minimum Rotated Rectangle

To measure width, a direction perpendicular to the contrail must be defined. Since a contrail can appear at any orientation in the image, as seen in Figure 4.4, this direction cannot be assumed and must be estimated from the polygon geometry. While the advected ADS-B flight track could in principle define a longitudinal axis, as has been explained above, ERA5 wind errors introduce positional uncertainty that grows with contrail age, meaning projected waypoints can be substantially displaced from the visible contrail over time. Instead, the longitudinal axis is estimated by computing the minimum rotated rectangle (MRR): the smallest-area rectangle of any orientation that fully encloses the polygon, using the Shapely library [49]. The MRR is shown as a dashed rectangle in the top right panel of Figure 4.6. Because a contrail is substantially longer than it is wide, the MRR is always elongated along the contrail direction, so its long axis defines the centerline along which cross-sectional widths are measured. Being derived entirely from the annotated polygon, the MRR is independent of advection accuracy. Its only limitation is that it requires the polygon to be longer than it is wide, a condition that holds for all observations in this dataset.

Width Measurement

At each position along the centerline, a cut line is drawn perpendicular to the longitudinal axis and intersected with the polygon. These cuts are shown as gray lines in the top right panel of Figure 4.6, with the maximum-width cut highlighted in red. The cut positions are taken at the projection of each polygon vertex onto the centerline: since the width function is piecewise linear between consecutive vertex projections, the maximum must occur at one of these positions. The length of each intersection in unit-height coordinates, w_{unit} , is converted to physical width by

$$w = w_{\text{unit}} \cdot h_{\text{rel}}. \quad (4.12)$$

The resulting width profile is shown in the bottom right panel of Figure 4.6. Each gray point corresponds to one vertex cut, and the red point marks the maximum. For this observation the width varies substantially along the contrail, reflecting the irregular polygon shape, which motivates the two summary metrics defined in the following section.

4.1.8 Width Metrics

Maximum Width

The maximum width is the largest single width measurement across all vertex-cut positions along the contrail:

$$w_{\text{max}} = \max_{k=1, \dots, N} w_k \quad (4.13)$$

This metric is considered the most relevant for satellite detectability. It is hypothesized that a contrail should be most visible at its widest point, and that one wide enough to influence the signal of at least one satellite pixel somewhere along its length would in principle be detectable, even if the rest is narrower. In the example of Figure 4.6, $w_{\text{max}} = 3.68$ km, occurring at the bulge visible near the center

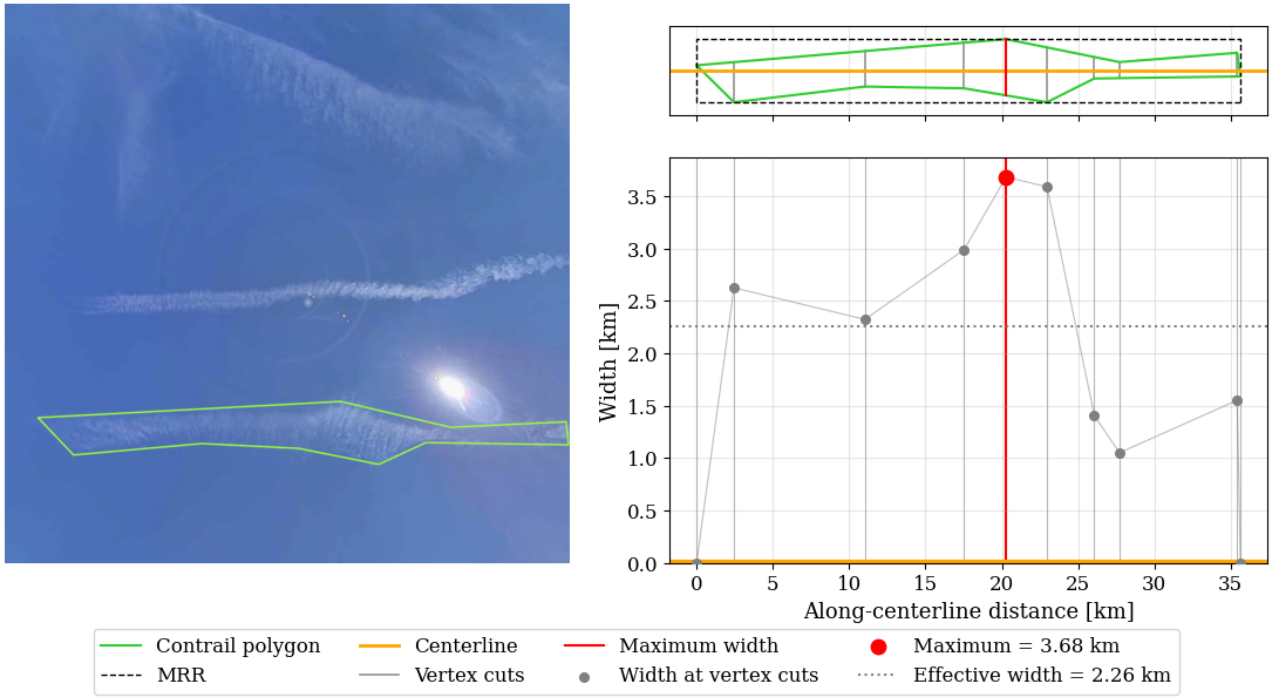


Figure 4.6: Width estimation for flight AAL2482, fifth observation, SGP, 09/12/2025 15:43 UTC. *Left:* ASI image with the annotated polygon overlaid in green. *Top right:* polygon in centerline coordinates; the dashed rectangle is the MRR, grey lines are perpendicular vertex cuts, and the red line marks the maximum-width cut. *Bottom right:* measured width at each vertex cut position, with the maximum (3.68 km) marked in red and the effective width (2.26 km) shown as a dotted line.

of the polygon. However, w_{\max} is sensitive to annotation of polygon corners: the sharp edges might locally inflate the measured width.

Effective Width

The effective width is defined as the contrail polygon area divided by the centerline length:

$$w_{\text{eff}} = \frac{A}{L}, \quad (4.14)$$

where A is the polygon area in m^2 and L is the centerline length in m. Geometrically, w_{eff} is the width of a rectangle with the same area and length as the contrail, and is therefore more robust to localized shape irregularities than the maximum. For the example in Figure 4.6, $w_{\text{eff}} = 2.26$ km, as shown in the dotted line of the bottom right panel. This is a 39% difference from maximum width. The area-weighted average is pulled below the peak because the width varies substantially along the contrail’s length, tapering at one end while the other end remains comparatively wide; the area average reflects the whole profile rather than the local maximum. w_{eff} is used frequently in contrail studies based on satellite imagery [7, 35], where the full contrail length is typically visible. However, from a ground camera, a contrail is rarely observed in its entirety. As it advects, one or both ends exit the frame, shortening the visible centerline length L over successive observations. This shrinks w_{eff} even if the contrail itself has not narrowed, making it sensitive to viewing geometry rather than physical change.

For this reason, w_{\max} is used as the primary width metric throughout this study. It corresponds directly to the widest point visible in the ASI, and since satellite detectability is assessed at that same

location, it is the most directly comparable measure for evaluating whether the contrail produces a detectable signal.

4.1.9 Sources of Uncertainty

Several sources of uncertainty affect the estimated contrail widths.

Altitude uncertainty

The limitations of the altitude estimate are discussed in Section 4.1.6: the assigned altitude is taken from the matched ADS-B waypoint at first observation and held fixed across all time steps, introducing uncertainty that grows as the contrail ages and spreads. Here, a further source of systematic offset is considered. Converting a width from pixel space to meters requires knowledge of the contrail altitude h , since $w_m = w_{\text{px}} \cdot h/f$. During the vortex phase immediately after formation, the lower part of a contrail is estimated to sink by 50–300 m depending on aircraft and atmospheric parameters [10], meaning the observed contrail may lie below the flight level recorded in the trajectory data. Accordingly, a downward offset of up to $\Delta h = 300$ m is assumed, modelled as a uniform distribution over the interval $[h - 300 \text{ m}, h]$.

Pixel annotation error

The contrail polygon boundaries are drawn manually by clicking vertices in the rectified all-sky image. This introduces a positional uncertainty on each clicked boundary point, arising from the finite pixel resolution of the image and the subjective judgment of where the contrail edge lies.

The width at a given cross-section is computed as the length of the intersection between a perpendicular cut line and the polygon, and the maximum across all cross-sections is taken as the contrail width. Although the polygon comprises many vertices, the width at any cross-section is locally governed by the two nearest boundary points, one on each side. The uncertainty in the measured width is therefore approximated as arising from two independent positional errors of standard deviation σ_p . Each click carries an independent positional error σ_p (in pixels), so by standard propagation of independent Gaussian uncertainties:

$$\sigma_{W_{\text{px}}} = \sqrt{\sigma_p^2 + \sigma_p^2} = \sqrt{2} \sigma_p. \quad (4.15)$$

Projected onto a horizontal plane at height h above the camera, one pixel corresponds to h/f meters, giving a physical width uncertainty of:

$$\sigma_W = \frac{\sqrt{2} \sigma_p h}{f}. \quad (4.16)$$

To estimate σ_p , the same contrail was annotated five times independently on a single image. The five resulting maximum width estimates had a sample standard deviation of $\sigma_W = 35.8$ m. Inverting Equation 4.16 with the known altitude and focal length $f = 260.1$ px yields σ_p as the standard deviation of a single edge-click position error in pixels. Since σ_p is calibrated from repeated annotations of the full polygon, it implicitly captures the collective effect of all vertex uncertainties on the final width estimate. This experiment was repeated on a second contrail, yielding values of 0.94 px and 1.12 px respectively. The larger value is adopted as a conservative estimate throughout this study.

Monte Carlo propagation

Because altitude uncertainty affects the pixel-to-meter conversion, the two error sources cannot be treated independently. Instead, they are propagated jointly using a Monte Carlo simulation with $n = 10,000$ samples per observation. For each sample, two quantities are perturbed simultaneously. First, a small random clicking error is added independently to each of the two boundary edges. Each error is drawn from a normal distribution with standard deviation $\sigma_p = 1.12$ px, representing the annotation uncertainty quantified earlier. The perturbed pixel width for sample i is:

$$\tilde{w}_{\text{px}}^{(i)} = w_{\text{px,nom}} + \varepsilon_{\text{left}}^{(i)} + \varepsilon_{\text{right}}^{(i)}, \quad \varepsilon_{\text{left}}, \varepsilon_{\text{right}} \sim \mathcal{N}(0, \sigma_p^2) \quad (4.17)$$

Second, the contrail altitude is drawn from a uniform distribution spanning the plausible range:

$$\tilde{h}^{(i)} = h - 300 \cdot U^{(i)}, \quad U^{(i)} \sim \mathcal{U}(0, 1) \quad (4.18)$$

Each perturbed sample then produces a physical width estimate:

$$\tilde{w}^{(i)} = \tilde{w}_{\text{px}}^{(i)} \cdot \frac{\tilde{h}^{(i)}}{f} \quad (4.19)$$

The 10,000 values of $\tilde{w}^{(i)}$ represent the range of width outcomes consistent with the annotation and altitude uncertainties. The 1σ uncertainty for each observation is taken as their standard deviation, and all width estimates are reported with $\pm 1\sigma$ uncertainty bands derived in this way.

4.2 Ambient Condition Estimates

Three ambient meteorological quantities are calculated for each contrail observation: The relative humidity over ice (RH_i), the Brunt–Väisälä frequency (N_{BV}), and the vertical wind shear (S).

Altitude–Pressure Conversion

ERA5 model-level data are provided on 137 hybrid sigma-pressure levels, with vertical spacing of roughly 200–300 m near cruise altitudes [11]. Levels 59–85 are used, covering altitudes of approximately 8.6–14 km and therefore bracketing all contrail altitudes in the dataset. Temperature T , specific humidity q , and wind components u and v are horizontally interpolated (longitude, latitude, time) to the contrail centroid location, retaining the full vertical profile. T and q are then vertically interpolated to the centroid altitude to compute RH_i and N_{BV} . Wind shear is computed as a finite difference of u and v between the two model levels that bracket the centroid altitude.

4.2.1 Relative Humidity with Respect to Ice

RH_i is computed from the interpolated temperature T (K), specific humidity q (kg kg^{-1}), and pressure p (Pa). The function calculates

$$RH_i = \frac{e}{e_{\text{si}}(T)}, \quad (4.20)$$

where the actual vapor pressure is approximated as $e = qp(R_v/R_d)$, with $R_v/R_d \approx 1.608$, the ratio of the specific gas constants of water vapor and dry air, and $e_{si}(T)$ is the saturation vapor pressure over ice evaluated using the Sonntag [50] formula. A value $RH_i > 1$ indicates ice-supersaturation, the thermodynamic condition required for contrails to persist [8].

4.2.2 Vertical Wind Shear

Vertical wind shear describes the rate of change of horizontal wind velocity with altitude. A high shear means the wind at the top and bottom of a layer differ in speed, direction, or both, which promotes the lateral spreading of contrails by spreading the ice crystal plume across the horizontal plane, which is expected to increase the apparent contrail width observed.

The wind components are interpolated to altitudes $z \pm 100$ m around the contrail altitude using four-dimensional linear interpolation, and the component-wise vertical shear is computed across this $\Delta z = 200$ m layer [51] as $S_u = \Delta u/\Delta z$ and $S_v = \Delta v/\Delta z$, following Schumann [10]. The total shear magnitude is then

$$S = \sqrt{S_u^2 + S_v^2}. \quad (4.21)$$

The cross-track component S_{\perp} is the projection of the shear vector onto the axis perpendicular to the contrail, since only this component displaces the contrail laterally and contributes to widening; along-track shear merely stretches the contrail lengthwise. The contrail bearing β , measured clockwise from North, is estimated from the long axis of the MRR of the polygon, giving

$$S_{\perp} = |S_u \cos \beta - S_v \sin \beta|. \quad (4.22)$$

This decomposition is illustrated in Figure 4.7. The total shear vector $\mathbf{S} = (S_u, S_v)$ is resolved into an along-track component S_{\parallel} , which stretches the contrail lengthwise without affecting its width, and a cross-track component S_{\perp} , which displaces ice crystals laterally and is therefore expected to drive width growth. Both S and S_{\perp} are expressed in units of $\text{m s}^{-1} \text{ km}^{-1}$.




Figure 4.7: Decomposition of the wind shear vector $\mathbf{S} = (S_u, S_v)$ relative to the contrail axis. The bearing β is measured clockwise from North. The cross-track component $S_{\perp} = |S_u \cos \beta - S_v \sin \beta|$ displaces ice crystals perpendicular to the contrail and drives lateral spreading; the along-track component S_{\parallel} stretches the contrail lengthwise and does not affect its width.

4.2.3 Atmospheric Stability

The Brunt-Väisälä frequency N_{BV} quantifies the static stability of the atmosphere at the contrail altitude, which describes the degree to which an air parcel resists vertical air motion. In a strongly stable atmosphere, vertical displacements are rapidly restored, confining a contrail to a thin horizontal layer; in a weakly stable atmosphere, the contrail can deepen vertically [52]. It is defined as

$$N_{BV} = \sqrt{\frac{g}{\theta} \frac{\partial \theta}{\partial z}}, \quad (4.23)$$

where $g = 9.81 \text{ m s}^{-2}$, θ is the potential temperature, and $\partial \theta / \partial z$ is its vertical gradient evaluated over a 200 m layer centered on the contrail altitude. The potential temperature is the temperature an air parcel would have if brought adiabatically to a reference pressure of 1000 hPa, given by $\theta = T(p_0/p)^{R/c_p}$, where T is the absolute temperature, p the ambient pressure, R is the specific gas constant of air and c_p is the specific heat capacity at a constant pressure.

4.3 Visibility in Satellite Imagery

Once a contrail polygon has been defined in the ground-based image, it must be projected into the satellite image to assess whether the contrail is visible there. However, in order to match the polygon in the ground camera image to its projection in the satellite, the parallax displacement must be corrected.

4.3.1 Parallax Correction

Because geostationary satellites view the Earth from an oblique angle, a contrail at altitude does not appear at its true geographic position in the satellite image. As shown in Figure 4.8, the satellite sees the contrail as though on the ground at position p' , rather than the true position p . The magnitude of this displacement depends on both the contrail altitude and how obliquely the satellite views the site.

To quantify the displacement, we first compute the angular separation γ at the Earth's center between the site and the satellite. Since the satellite is located above the equator, this reduces to a simple function of the site latitude φ and the longitude difference $\Delta\lambda$ between the site and the satellite:

$$\cos \gamma = \cos(\Delta\lambda) \cos \varphi. \quad (4.24)$$

The slant range s from the satellite to the contrail follows from the law of cosines in the Earth-center–contrail–satellite triangle:

$$s = r_{\text{sat}} \sqrt{1 + \left(\frac{R}{r_{\text{sat}}}\right)^2 - 2 \frac{R}{r_{\text{sat}}} \cos \gamma}, \quad (4.25)$$

where $R = R_E + h_{\text{contrail}}$ is the distance from the Earth's centre to the contrail. The zenith angle θ_z then follows from the sine rule:

$$\sin \theta_z = \frac{r_{\text{sat}}}{s} \sin \gamma, \quad (4.26)$$




Figure 4.8: Geometry of a contrails parallax shift. A contrail at altitude h_{contrail} above site p appears at position p' in the satellite image, displaced horizontally by Δd . The zenith angle θ_z is measured at the contrail between the local vertical and the line of sight s to the satellite.

where $r_{\text{sat}} = R_E + H_{\text{sat}}$ is the orbital radius. The elevation angle is $\theta_e = 90 - \theta_z$. The horizontal displacement is then simply

$$\Delta d = \frac{h_{\text{contrail}}}{\tan \theta_e}. \quad (4.27)$$

The direction of the displacement is radially away from the sub-satellite point: at SGP the satellite lies to the south-east, so contrails appear shifted north-west; at CAB the satellite is almost due south, so the shift is mostly northward.

4.3.2 Brightness Temperature Calculation

The parallax-corrected contrail polygon is used to identify the brightness temperature difference signal of the contrail. The BTD signal is recorded for each pixel inside the polygon, and the mean and standard deviation are computed:

$$\mu_{\text{BTD, polygon}}, \quad \sigma_{\text{BTD, polygon}} \quad (4.28)$$

The standard deviation within the polygon reflects the spatial variability of the contrail signal: a high value may indicate that the contrail is not uniformly visible and the polygon might not be capturing the contrail properly.

To estimate the local background, a region around the polygon is defined. The MRR is scaled only in the cross-track direction, not along the contrail axis, as that would risk including more of the

contrail itself in the background region. Pixels within the inner rectangle and within the polygon itself are excluded. The mean and standard deviation of BTD in this background region are:

$$\mu_{BTD, bg}, \quad \sigma_{BTD, bg} \quad (4.29)$$

A high $\sigma_{BTD, bg}$ indicates a noisy background – for example due to nearby natural cirrus – which reduces confidence in the contrast estimate.

The contrast is then calculated by:

$$BT_{\text{contrast}} = \mu_{BTD, \text{polygon}} - \mu_{BTD, bg} \quad (4.30)$$

A positive value indicates the contrail pixels are warmer in BTD than their surroundings, which is the expected signal of a contrail at cruise altitude. An example is shown in Figure 4.9, where the contrail is clearly distinguishable as a dark blue linear feature.

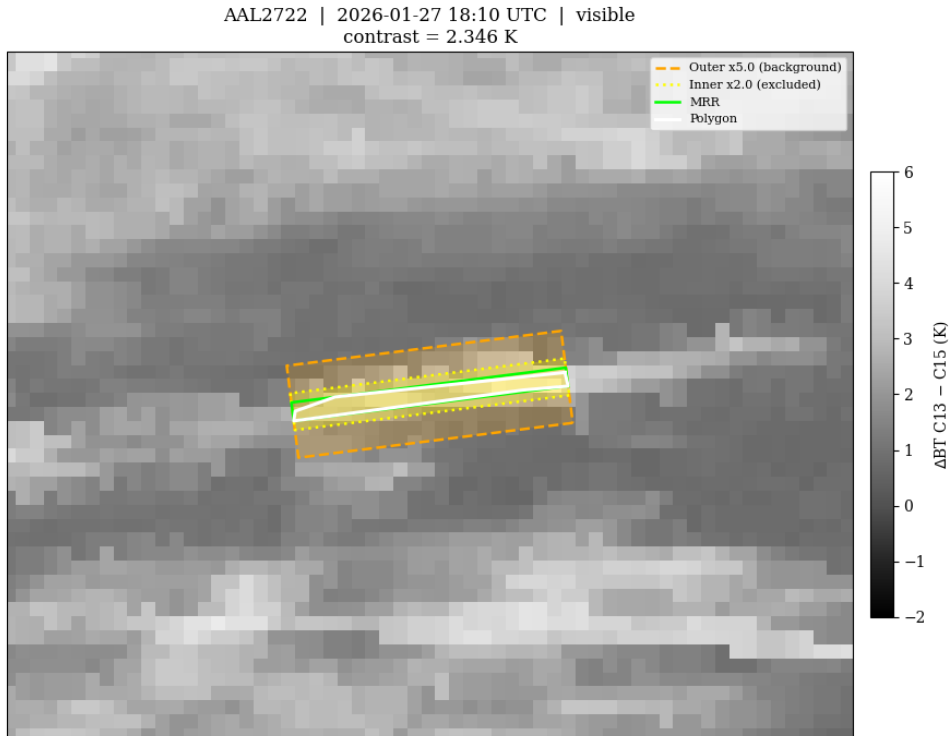


Figure 4.9: Example BTD field (C13–C15) for contrail produced by flight AAL2722. The white outline shows the parallax-corrected contrail polygon. The green rectangle is the MRR; the orange dashed and yellow dotted rectangles mark the outer ($\times 5$) and inner ($\times 2$) background boundaries respectively.

Choice of background region size

The outer scale factor s_{outer} determines how far from the contrail the background is sampled. In order to determine the best value for this factor, the contrast was computed for outer scale factors ranging from 3 to 8 across six contrails with varying signal background standard deviation. For contrails with a low $\sigma_{BTD, bg}$, meaning their background is uniform, most likely clear sky, the contrast stabilizes around $s_{\text{outer}} = 5$, beyond which it begins to decline as the background region grows large enough to sample atmospheric variability further from the contrail. For contrails with a high $\sigma_{BTD, bg}$, the contrast continues to grow with scale factor, reflecting an inhomogeneous background that is not well captured by a small region. A value of $s_{\text{outer}} = 5$ was therefore adopted as a compromise: it yields a stable estimate for low-noise backgrounds without extending so far that the background no longer represents local atmospheric conditions near the contrail.

4.3.3 Effective Resolution at the Sites

Before assessing contrail visibility in the satellite imagery, it is necessary to establish the effective ground resolution at each site. This is visualized for simplicity in Figure 4.10. The nominal pixel footprint of a satellite instrument is defined at nadir, but degrades for off-nadir viewing. A fixed instrument angular resolution δ maps to a footprint $d_{\text{nadir}} \approx H\delta$ at the sub-satellite point, where H is the orbital altitude. At a ground point where the satellite appears at zenith angle θ_z , the slant range increases to $H/\cos\theta_z$, and the projected footprint scales accordingly to:

$$d_{\text{eff}} = \frac{d_{\text{nadir}}}{\cos\theta_z} \quad (4.31)$$

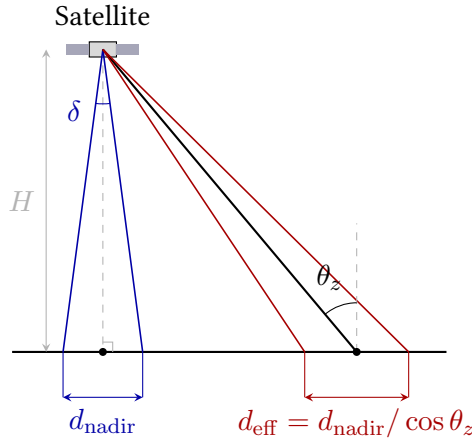


Figure 4.10: Effective pixel footprint geometry. A fixed instrument angular resolution δ maps to $d_{\text{nadir}} \approx H\delta$ at the sub-satellite point (blue). At an off-nadir site viewed at zenith angle θ_z , the same angular aperture projects to $d_{\text{eff}} = d_{\text{nadir}}/\cos\theta_z$ (red).

The zenith angle θ_z at each site follows from the same orbital geometry as the parallax correction in subsection 4.3.1, setting $h_{\text{contrail}} = 0$ so that the line of sight traces to the surface. The resulting viewing geometry and effective footprints are summarized in Table 4.2.

Table 4.2: Satellite viewing geometry and effective pixel footprint at each ground site. Nadir resolutions are 2 km for the GOES-East Cirrus RGB product and 1 km for the MTG-I1 Cloud Type RGB product.

	SGP	CAB
Satellite	GOES-East	MTG-I1
Sub-satellite lon.	75.2 W	0.0 E
Central angle γ	42.1	52.1
Slant range s	37,675 km	38,575 km
Zenith angle θ_z	48.6	59.6
Nadir resolution	2 km	1 km
Effective resolution	3.02 km	1.98 km

Substituting the nadir resolutions into Equation 4.31 gives 3.02 km at SGP and 1.98 km at CAB. Although MTG-I1 begins with a factor-of-two nadir advantage over ABI, the larger zenith angle at CAB partially offsets this gain. The ratio of effective footprints is

$$\frac{d_{\text{eff,SGP}}}{d_{\text{eff,CAB}}} = \frac{d_{\text{nadir,GOES}}}{d_{\text{nadir,MTG}}} \cdot \frac{\cos \theta_{z,\text{CAB}}}{\cos \theta_{z,\text{SGP}}} = \frac{2}{1} \cdot \frac{\cos 59.6}{\cos 48.6} \approx 1.53, \quad (4.32)$$

so the nominal $2\times$ advantage is reduced to roughly $1.5\times$ in effective footprint. MTG-I1 therefore retains a resolution advantage, but the two sites are more comparable in practice than their nominal specifications suggest.

4.3.4 Background Condition Labels

To enable a later comparison of contrail observability in satellite images under differing background conditions, each observation date was assigned a sky condition label. These labels were assigned manually for each observation date by inspecting the Ash and Cloud Type RGB composites in each site alongside the ASI imagery. An observation was classified as "cirrus" if the RGB showed a clearly populated cirrus layer over the site, and as "clear" otherwise. The threshold is necessarily subjective, but the classification is intended to reflect whether background conditions were visually cirrus-heavy enough to plausibly interfere with contrail detection. Figure 4.11 shows a representative example of each condition at SGP.

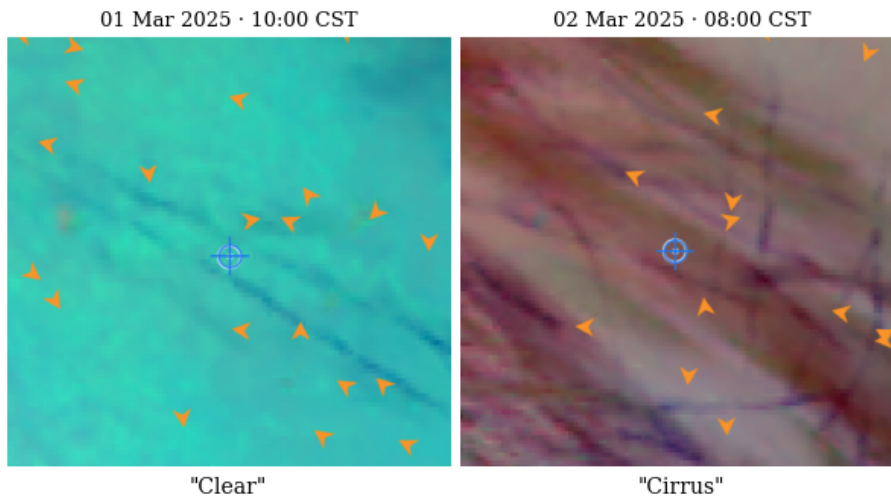


Figure 4.11: Ash composites from GOES-East ABI illustrating the two classified sky conditions. *Left*: clear sky, 1 March 2025 10:00 CST. *Right*: cirrus, 2 March 2025 08:00 CST.

Results

5.1 Dataset Overview

The final dataset consists of 148 individually tracked contrails: 100 observed at SGP across 23 days between January 2025 and February 2026, and 48 at CAB across 6 days between March and November 2025. At SGP, 49 contrails were classified as visible in GOES-East imagery at some point during the time they were observable in the ground-based imagery. At CAB, 24 contrails became visible in MTG imagery. In total, the dataset comprises 740 individual polygon observations. Each contrail was observed across multiple timesteps, yielding a minimum of two but up to thirteen polygon annotations per contrail depending on its visibility window in the ASI FOV. The total duration for which individual contrails were tracked ranged from 8.7 to 85 minutes, with a mean of 29 minutes. Table 5.1 summarizes the key properties of each site subset.

Table 5.1: Summary of the contrail dataset collected at each site.

	SGP	CAB
Satellite instrument	GOES-East ABI	MTG-I1 FCI
Observation period	Jan 2025 – Feb 2026	Mar – Nov 2025
Days with observations	21	6
Total contrails	100	48
Visible	49	24
Not visible	51	24
Total polygon observations	480	260

Temporal Distribution

Figure 5.1 shows the monthly count of contrails observed at SGP across the days analyzed. The distribution peaks in the winter and spring months, with a broad trough from May through October. Zero counts in May, June, and October do not reflect an absence of contrail activity, they reflect periods during which no qualifying Google Contrails events, as described in subsection 3.1.1, were

identified and manually reviewed. The dataset therefore represents a targeted sample of events selected for satellite contrail detection analysis, rather than a representative sample of all contrail occurrences above the site. Despite these gaps, the observed distribution is consistent with the broader seasonal pattern reported in previous studies, with higher occurrence in the winter and early spring months, a decline from May onward, and a minimum in midsummer before recovering toward the end of the year [14, 43, 26]. As the CAB dataset does not span a full calendar year, no equivalent monthly distribution is shown for that site.




Figure 5.1: Monthly count of individually tracked contrails at the SGP site in 2025. Zero counts in May, June, and October reflect gaps in manual event selection rather than zero contrail occurrence.

Figure 5.2 shows the diurnal distribution of contrail observations at both sites. Observations across the 21 analyzed days at SGP are concentrated in the morning hours, peaking around 8:00 AM local time and declining through the afternoon. This broadly reflects the diurnal pattern in contrail coverage reported by Meijer et al. [26] for central United States, where coverage peaks in the early morning and decreases through midday. At CAB, the distribution peaks around noon. The apparent morning peak at CAB at 7:00 AM local time is an artifact of the limited size of the dataset: seven of the 48 CAB contrails were captured on a single morning (6 September 2025) and should not be interpreted as a systematic diurnal signal. With only 6 observation days, the CAB hourly distribution is sensitive to individual days and does not support conclusions about diurnal variability.

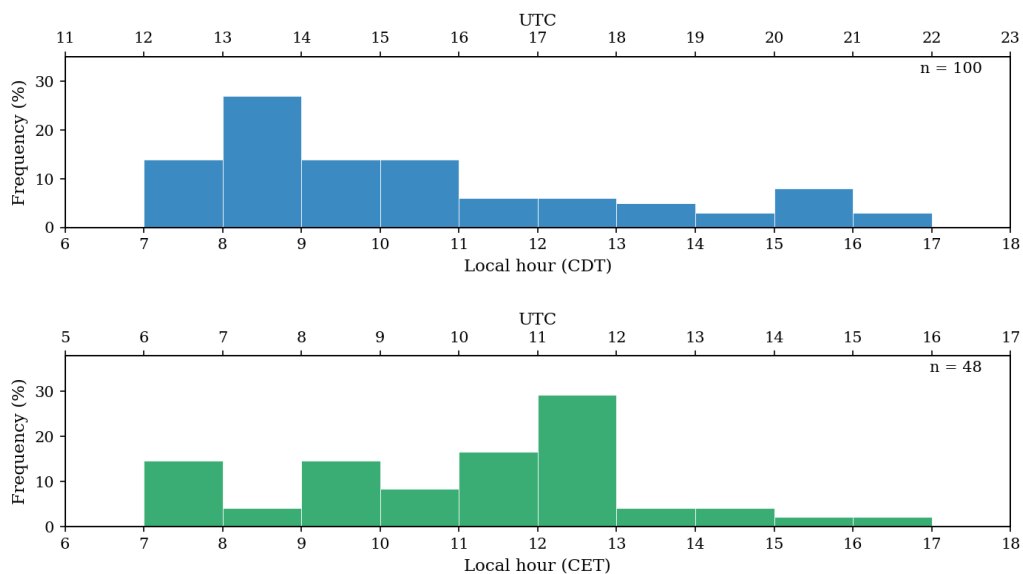


Figure 5.2: Hourly distribution of contrail observations at SGP and CAB.

5.1.1 Collected Variables

For each polygon observation, a set of variables is calculated and recorded spanning four categories: contrail geometry, flight data, atmospheric conditions at flight level, and satellite-derived signal. The calculation of these properties is detailed in chapter 4. Table 5.2 lists the full set of variables available for every observation.

Table 5.2: Variables per polygon observation.

Variable	Unit	Source
<i>Contrail geometry</i>		
Maximum width	m	ASI polygon
Effective width	m	ASI polygon
Width uncertainty	m	ASI polygon
Lateral growth rate	km min ⁻¹	ASI polygon
<i>Flight data</i>		
Callsign	–	ADS-B
Contrail age	min	ADS-B
Flight altitude	m	ADS-B
Geographic position	°	ASI polygon
<i>Atmospheric conditions</i>		
Relative humidity over ice	–	ERA5 model levels
Brunt-Väisälä frequency	s ⁻¹	ERA5 model levels
Vertical wind shear	m s ⁻¹ km ⁻¹	ERA5 model levels
Sky condition	–	Manual label
<i>Satellite image signal</i>		
BTD polygon	K	GOES-East ABI / MTG-I1 FCI
BTD background	K	GOES-East ABI / MTG-I1 FCI
BTD contrast	K	GOES-East ABI / MTG-I1 FCI

The variables recorded represent what can be derived about a contrail from the available ground-based, satellite, flight and weather data, and will be used to answer this study’s two research questions:

- Which combination of contrail properties, background conditions, and instrument characteristics determines whether a contrail is observable in a given satellite image?
- How long does it take for a contrail to become observable in satellite imagery, and how does this delay depend on atmospheric and background conditions?

This chapter will first explore what determines observability in section 5.2, and then what determines the delay in observability in section 5.3.

5.2 What Determines Observability?

This section addresses the first research question by examining the contrail properties, background conditions and instrument characteristics that determine whether a contrail is detectable in geostationary satellite imagery. First, contrail width is examined in subsection 5.2.2, then age at first satellite detection in subsection subsection 5.2.3, and finally a brightness temperature analysis of

each contrail polygon to assess whether an objective visibility measure is consistent with the manual labels is presented in subsection 5.2.5.

5.2.1 Visibility Label

Before presenting results, it is important to define what is meant by "observability" in this study. A contrail is classified as *visible* in a satellite image if it can be identified as a coherent linear feature in the Ash or Cloud Type RGB at the expected location. This assessment is made manually: each contrail was first identified and annotated in the ASI image on VOCASION, which then projects the contrail polygon to its location in the satellite image, as seen in Figure 5.3, and the annotator checks whether a matching linear feature is present within or near the polygon boundary. Successive satellite and ground images are scrolled through to confirm that the feature is evolving consistently with what is observed from the ground at the same time. If no corresponding feature can be identified, the observation is labeled *not visible*.

This approach nonetheless introduces subjectivity. There could be contrails close together, overlapping, or heavy natural cirrus cloud cover, and different labelers might reach different conclusions for whether a contrail is visible or not in some difficult cases. Knowing the expected contrail location also introduces a risk of confirmation bias: a signal inside a polygon which is surrounded by cirrus may be labeled visible simply because the labeler sees dark blue pixels in that area and knows where to look, rather than because the contrail is genuinely distinguishable from the surrounding cirrus. An example of a difficult case in an Ash RGB image can be seen in Figure 5.3. The screenshot is taken from VOCASION on a day with a lot of natural cirrus and contrail formation. The contrail that has been annotated is wide enough to be expected to be observable in the satellite image, however due to the surrounding cirrus, it is difficult to resolve it. This is a case which was labeled *not visible* in the contrail dataset, however a different labeler might disagree. Six more case studies of contrails in Ash and Cloud Type RGB images are shown in this chapter, and can be taken as examples of what different sky conditions might look like on different days above the two sites.

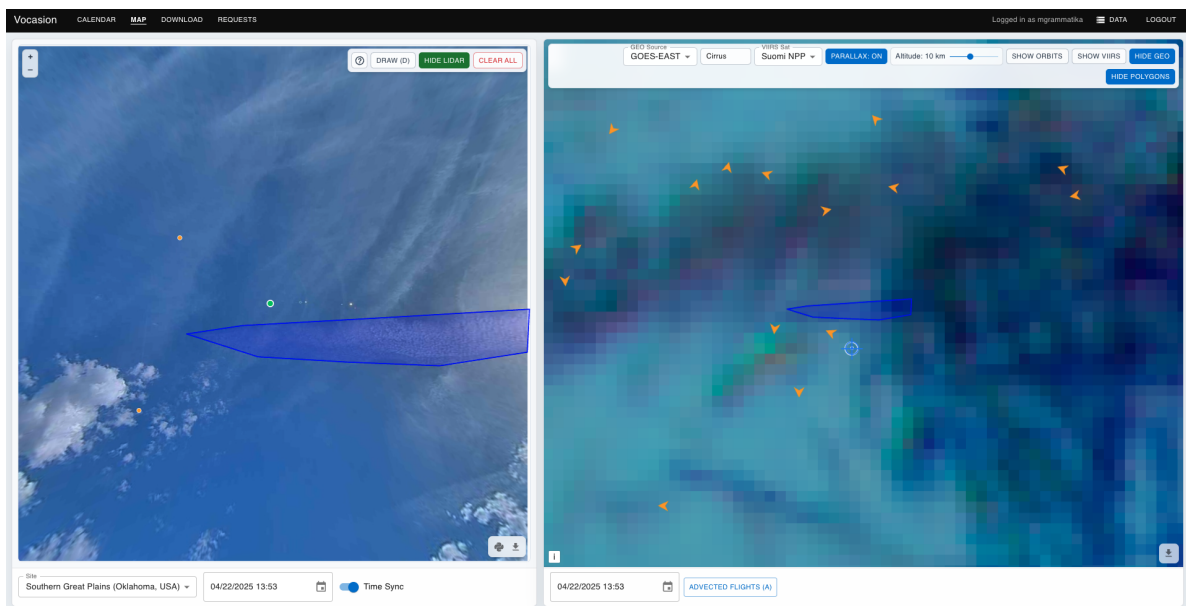


Figure 5.3: Screenshot from VOCASION for a contrail observation on 22 April 2025 at 13:53 UTC. The annotated contrail polygon is projected onto the Ash image, but the surrounding natural cirrus makes it difficult to determine whether the contrail is genuinely distinguishable from the background. This observation was labeled *not visible*.

5.2.2 Contrail Width

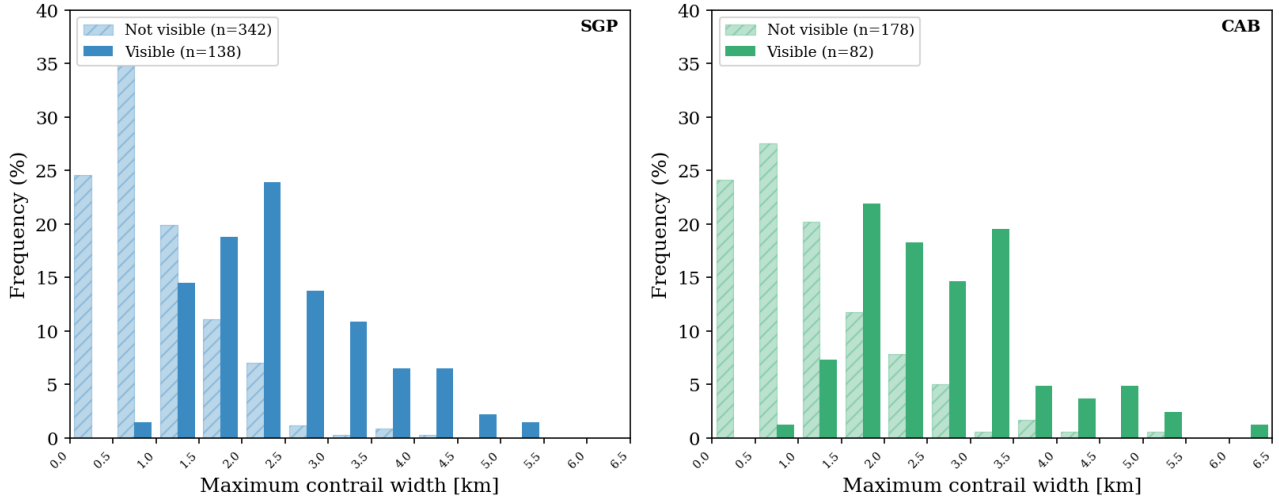


Figure 5.4: Distribution of maximum contrail width for visible (opaque) and not-visible (hatched) satellite observations at SGP and CAB.

Figure 5.4 shows the distribution of maximum contrail width for all 740 observations, separated by visibility label. Each observation is treated as an independent data point, meaning that a single contrail contributes multiple entries (one per annotated time step) depending on how long it remained in the FOV. At both sites, there is a separation between visible and not-visible distributions. The visible distribution extends to substantially larger widths, peaking at around 2 km widths and extending until 6.5 km widths, with non-visible observations becoming increasingly rare beyond 3 km. This reflects the fact that once a contrail is first detected in satellite imagery, it continues to be annotated and therefore contributes wider measurements to the visible distribution as it grows. Non-visible observations are heavily concentrated below 1.5 km, peaking below 1 km. This is partly expected, as narrower contrails are less likely to meet the detection threshold of satellite imagery. However, the non-visible distribution also reflects a sampling constraint: observations end when the contrail leaves the ASI field of view, meaning contrails that had not yet reached satellite visibility at that point are classified as never-visible regardless of whether they subsequently became detectable. As this study does not classify whether a contrail dissipated or advected out of the field of view, contrail lifetime and persistence beyond the observation window cannot be quantified.

Figure 5.5 shows the distribution of maximum contrail width at the moment of first satellite detection. Three contrails from each site have been excluded because they were already visible in the satellite image as they advected into the ASI FOV; their width at first observation is therefore not representative. The minimum, maximum and mean widths at each site are reported in Table 5.3. At both sites, a subset of contrails already become visible at widths under 1 km. This is sub-pixel at the satellite imager resolutions used here, which is consistent with Mannstein, Brömser, and Bugliaro [6], who showed that contrails as narrow as 0.5–0.9 km can be detected in MSG SEVIRI imagery (3 km resolution), confirming that sub-pixel contrail widths can be detected.

The rare cases where contrails remain invisible in the satellite image at widths above 3 km are due to either heavy cirrus cover on the day, as in the case of AAL1985 (Figure 5.3), or a contrail that appears faint and low-contrast even at large widths. The contrails which do become visible but only above 3 km width are also due to either heavy cirrus background making visibility at narrower widths difficult, or a gradual increase in opacity over time, as in the case of the contrail produced by flight JBU577 (Figure 5.7), which is analyzed further in this chapter.

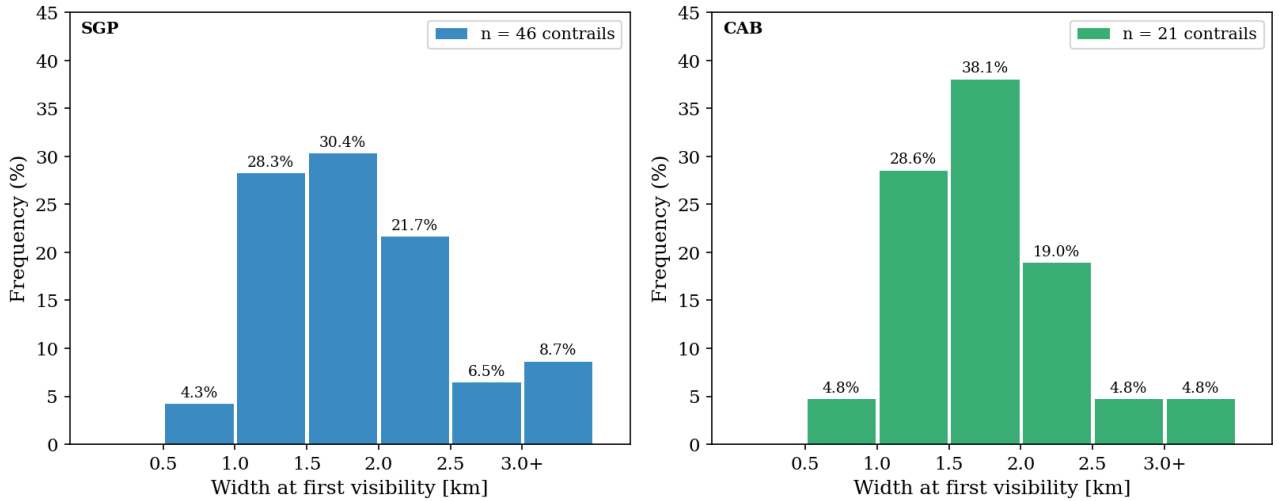


Figure 5.5: Maximum contrail width at the moment of first satellite detection at SGP and CAB sites.

Site	Min [km]	Max [km]	Mean [km]
SGP	0.724	3.436	1.86
CAB	0.842	3.030	1.79

Table 5.3: Contrail width statistics at first satellite detection.

Despite the difference in nominal spatial resolution between GOES-East ABI (2 km) and MTG-I1 FCI (1 km), the two width distributions in both figures are very similar, with only a slightly higher concentration of contrails visible at smaller widths at CAB. This could be explained by the satellite viewing zenith angle, the angle between the satellite’s line of sight and the vertical at the observed location, which is 48.6° at SGP and 59.6° at CAB. The larger zenith angle at CAB means the satellite views the site more obliquely, stretching the pixel footprint, and partially offsetting the resolution advantage of MTG-I1. Furthermore, in their study observing simulated contrails at different imager resolutions, Driver, Stettler, and Gryspeerdt [35] found that moving from a 2 km to a 0.5 km imager increased contrail observability by only 7 percentage points and suggested that the majority of undetectable contrails are too optically thin rather than too narrow. This could also explain the similarity in visible width distributions despite the resolution difference.

5.2.3 Age at First Observation

Figure 5.6 shows the distribution of contrail age at the moment of first satellite detection. At SGP, contrails are first detected at a mean age of 21.3 minutes, with the majority falling in the 15–20 minute bin. At CAB the mean detection age is 19.0 minutes, slightly lower than at SGP, and while the distributions are broadly similar, both peaking in the 15-20 minute bin, the entire distribution of CAB is shifted to slightly earlier detections. Most notably, approximately 10% of CAB contrails are first detected 5–10 minutes after formation, something not observed at SGP. However, given the smaller CAB sample size, it is unclear whether this reflects a genuine difference or a sampling artifact. Notably, the widths at first satellite detection are similar between sites (Table 5.3), which complicates a resolution-based explanation: if the FCI’s higher resolution were the primary driver of earlier detections, one would expect CAB contrails to be detected at smaller widths rather than simply younger ages. The earlier detections at CAB may therefore reflect higher contrail spreading rates at that site, resulting in contrails reaching detectable widths more quickly, though the dataset is too limited to draw firm conclusions.

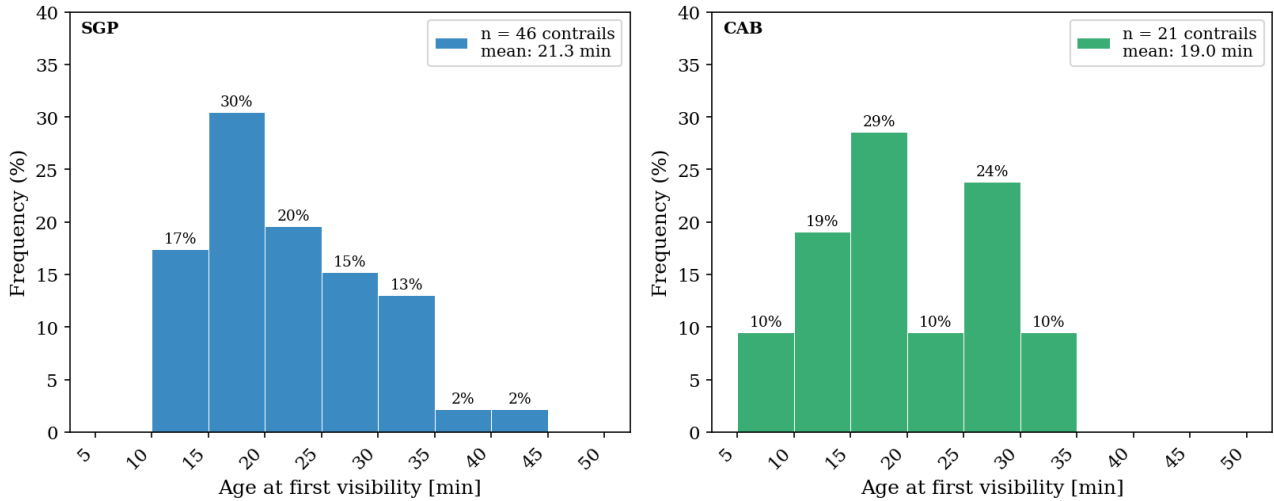


Figure 5.6: Age at first satellite detection at SGP and CAB.

The wide spread in age at first satellite observation could be a reflection of several factors: the spreading rate of a contrail, ambient conditions surrounding it and background sky conditions all influence when a contrail becomes detectable, and each of these will be examined in this chapter.

One factor this dataset cannot quantify is optical depth. The two detections beyond 40 minutes at SGP are cases where contrails exceed 2 km in width under clear sky conditions, yet only become detectable late in their lifetime. In the ground imagery, these contrails visibly transition from near-transparent to opaque over time; one hypothesis is that the contrail gradually increases in optical depth as ice crystals grow in the ice-supersaturated environment, eventually crossing the satellite detection threshold [4, 52]. One such example is examined in detail in Figure 5.7.

These detection ages are broadly consistent with existing literature. Previous studies using GOES-16 ABI imagery such as Gryspeerd et al. [34], who analyzed 64 000 contrail segments over the Western North Atlantic using several infrared bands, found a mean age to detection of 17 minutes, with 77% of contrails becoming visible in the first 20 minutes of formation. Similarly, Chevallier et al. [53] reported a mean detection age of 27 minutes after formation, during an 8-hour test case of automated contrail detections in Ash RGB images above the United States. Geraedts et al. [32] analyzed 1.6 million flight segments and found a mean time until first observation of 41 minutes, with most contrails first appearing about half an hour after formation.

Example of delayed detection

The contrail produced by flight JBU577, observed at SGP on 1 January 2026 at 09:42 local time, is first detected in satellite imagery 40 minutes after formation. At first ground observation, shown in the left panels, the contrail advects into the camera frame 30 minutes after formation at a width of 1.768 km. This width would generally be expected to be sufficient for satellite detectability, yet the contrail does not appear in the satellite image, suggesting that width alone was not the limiting factor. Five minutes later, shown in the center panels, the contrail appears visibly more opaque in the ground camera and has reached a width of 2.1 km. Ten minutes after first ground observation, in the rightmost panels, it has become sufficiently opaque to appear in the Ash RGB. This progressive increase in opacity suggests a gradual build-up of optical depth as a possible explanation for the delayed detection, though this cannot be confirmed as optical depth is not directly measured in this study. With a first detection width of 2.5 km and an age of 40 minutes, this contrail represents some of the largest values in this dataset. Formed under clear sky conditions, this case illustrates that satellite detectability may depend not only on geometric width but also on optical properties that

are not captured by the instruments used in this study.

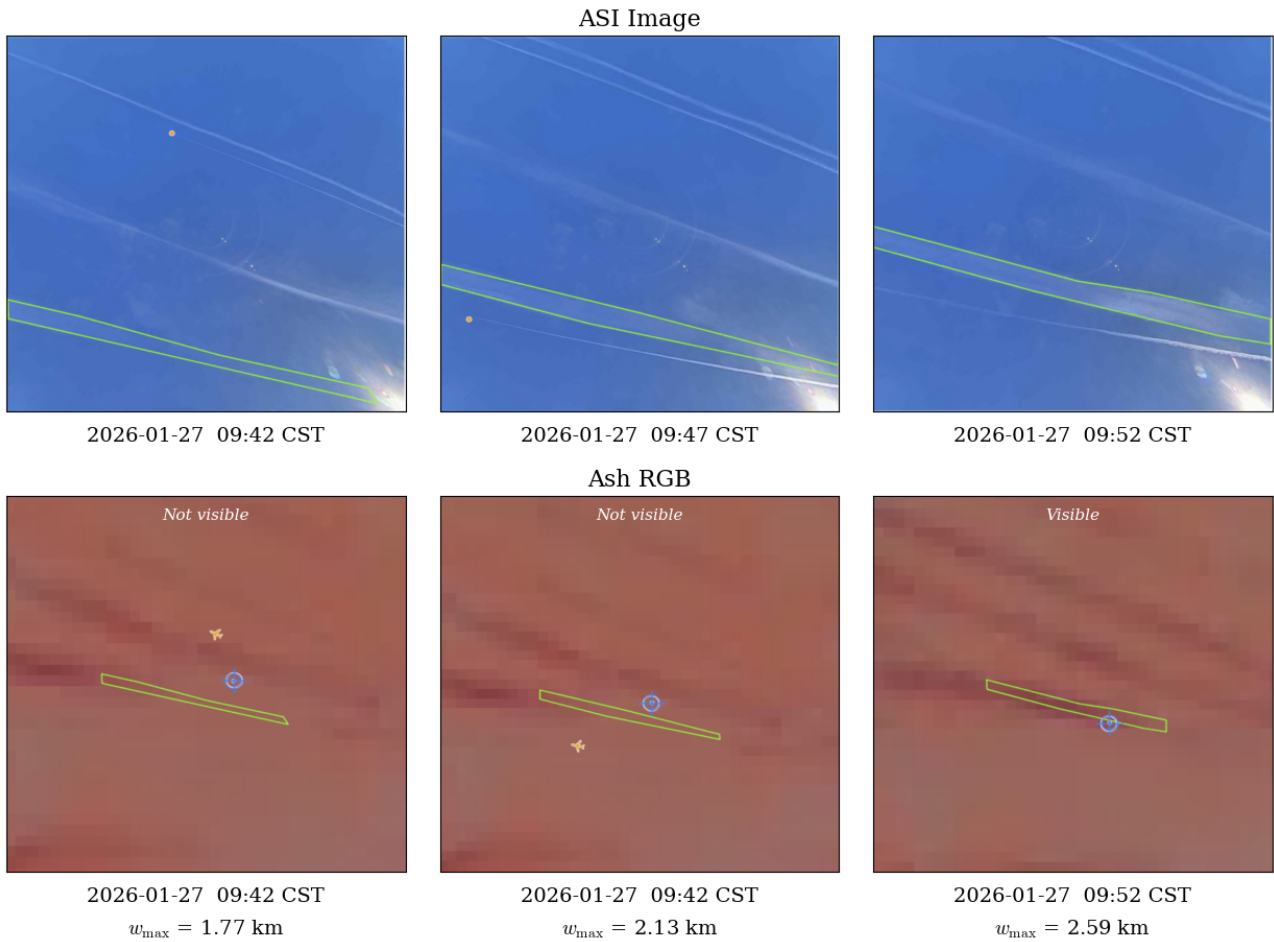


Figure 5.7: Contrail formed by flight JBU577 at 09:42 CST, shown at 5-minute timesteps, with the annotated polygon and the maximum width measured denoted at each timestep. In the Ash RGB images, the blue marker denotes the SGP site location, and the timestamp indicates the scan time of the satellite image.

5.2.4 Age-width relationship

Figure 5.8 synthesizes the joint dependence of satellite visibility on contrail age and maximum width, across all 480 observations at SGP. Each cell in the grid is colored by the fraction of observations in that bin classified as visible. The marginal histograms show the distributions of age and width at first satellite observation, identical to the histograms shown in the preceding sections. For CAB, with 260 observations, most cells in the equivalent heatmap are either empty or contain too few observations for the per-cell fractions to be interesting, and the plot is therefore not included.

Moving toward the top-right corner of the grid, where contrails are both older and wider, the fraction of visible observations increases. Older, wider contrails are more likely to be detectable than narrower, younger ones, and the transition toward higher visibility is most apparent at widths above 2 km and ages beyond 20 minutes. However, observations also become increasingly sparse in this top-right region, which is consistent with the study design. Most contrails form above the camera, advecting out within the first 30 minutes of their lifetime. As a result, the majority of observations are concentrated at younger ages and narrower widths, and contrails reaching large widths at old ages are more rare. These cells contain few observations and carry less statistical weight. Furthermore, the dataset is not sampled randomly: the aim was to have a roughly equal

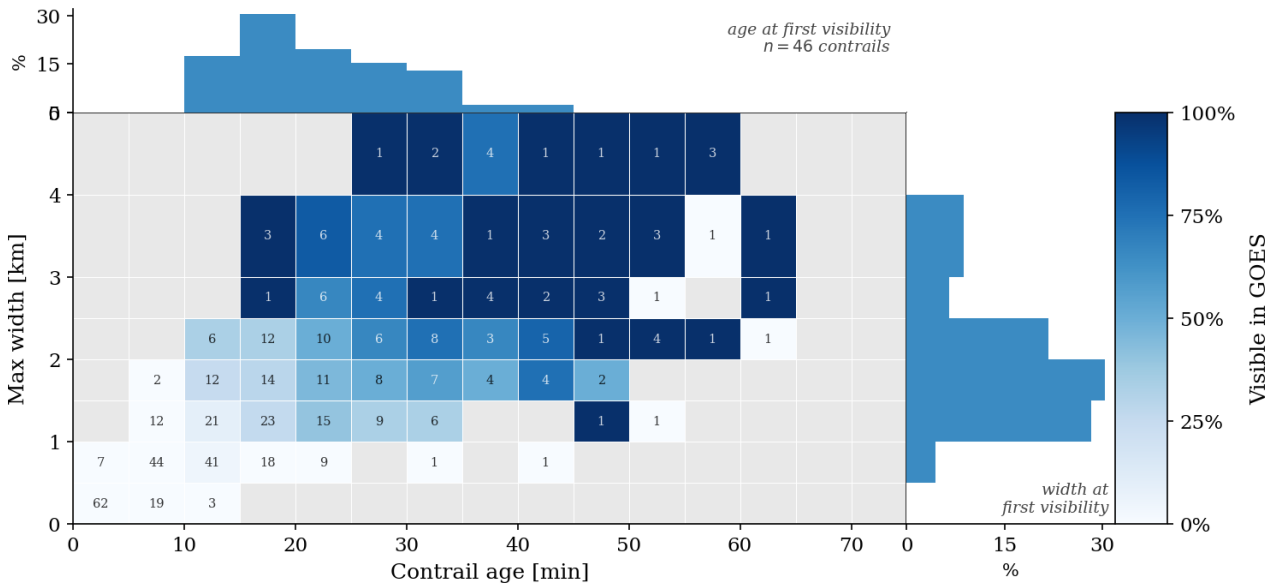


Figure 5.8: Satellite visibility fraction as a joint function of contrail age and maximum width at SGP. Cell color encodes the fraction of observations in that (age, width) bin that are visible in GOES; numbers give the observation count per cell. Marginal histograms show the distribution of age (top) and width (right) at first visibility.

number of visible and non-visible contrails, so the fraction of visible against visible observations reflects deliberate sampling choices rather than the true rate of satellite detection at a given age and width. This means that this figure is a visual insight of the general pattern of observability behavior of contrails in this dataset, but should not be interpreted as a general quantitative probability of observation.

5.2.5 Brightness Temperature Difference at First Observation

For an objective measure of contrail observability in this dataset, the BTM signal within each annotated contrail polygon projected onto the satellite image is analyzed.

SGP site: GOES-East ABI

At the SGP site, channels C13 (10.3 μm) and C15 (12.3 μm) of the GOES-East ABI were downloaded for each observation timestamp and cropped to the SGP site. For each parallax-corrected polygon annotation, the BTM contrast is computed as the mean brightness temperature difference within the polygon minus the mean of the surrounding background region. Figure 5.9 shows five distributions of this contrast value. The first two group all 480 polygon observations by visibility label, treating each timestep as an independent sample. The third and fourth distributions show, per contrail, only the first and last observation at which the contrail was classified as visible, capturing the contrast at the moment of satellite detection and at the end of the visible window respectively. The fifth distribution is a control group: 106 polygons were manually drawn at random positions over sky in the ASI imagery, at varying cloudy and clear conditions, with no associated contrail, and processed to obtain their BTM contrast values. This control group represents the background noise level of the contrast metric in the absence of any contrail signal.

The BTM contrast distribution can be used to assess the subjectivity of the observability labels created by means of manual classification. Every observation labeled visible has a positive BTM

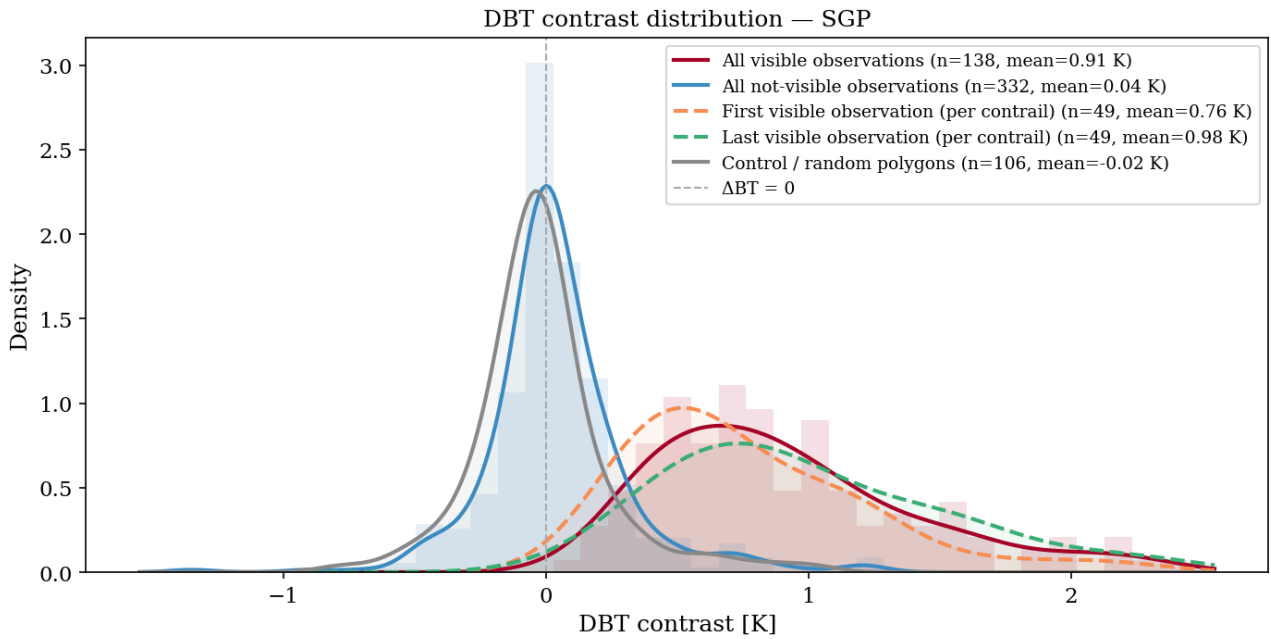


Figure 5.9: BTB distribution at SGP for visible, non-visible, and control observations. The visible distribution is shifted to positive values, validating the subjective visibility labels.




Figure 5.10: BTB contrast versus maximum contrail width at first visibility at SGP

contrast, which can be interpreted as having a contrail present. The visible distribution has a mean value of 0.91K. The non visible BTB contrast distribution is centered around 0 K, with a mean value of 0.04 K, meaning little difference between the polygon and its background, and matches almost exactly the distribution of a control group of randomly positioned polygons drawn over a wide variety of background sky conditions. The close agreement between the not-visible and control distributions validates the manual labeling in this dataset.

A secondary trend is visible across the observation lifetime of each contrail: the BTD contrast mean increases from first to last visible observations, as can be seen by the dotted orange and green distributions respectively. This is consistent with the contrail widening and becoming optically thicker over time, both of which amplify the contrail signature contrast against the background. An example of the progression from weak to strong signal within a single contrail's lifetime can be seen in the case study in Figure 5.11.

Figure 5.10 shows the relationship between BTD contrast and maximum width at first visibility. Every first visible observation has a contrast above 0.2K. This 0.2K minimum contrast value is identical to the one found by Mannstein, Brömser, and Bugliaro [6] when analyzing the BTD of contrails found in AVHRR (1.1 km resolution) Channel 4 (10.77 μm) minus Channel 5 (12 μm). Mannstein, Brömser, and Bugliaro [6] also found an average contrast value of 0.8K with a maximum value of 2-3K, overall matching the results of this analysis. The scatter shows a positive trend: wider contrails tend to produce stronger signals.

The contrails detected below 1 km width are particularly interesting, because despite their narrow geometry they still produce a contrast above the detection threshold. The two case studies that follow examine specific contrails in detail: Figure 5.11 shows a clear example of width and BTD signal growing together over time, and Figure 5.12 examines the sub-1 km width contrail which is visible on both Ash RGB and BTD images.

Case study: temporal evolution – SIA7429




Figure 5.11: Temporal evolution of the SIA7429 contrail in ASI (top) and BTD (bottom) over 30 minutes. The green dot marks the SGP ASI location for spatial reference.

Figure 5.11 illustrates the temporal evolution of contrail BTD signal over 30 minutes for the contrail produced by flight SIA7429. The upper panels show the ASI images, in which the contrail clearly forms and broadens. The lower panels show the corresponding BTD signal at each annotated polygon. The growing contrail produces an increasingly distinct thermal signature, consistent with the trends discussed above.

Case study: sub-1 km width at first visibility – SKW4822 The contrail produced by flight SKW4822 is first detected in the GOES-East Ash image at a width of 0.73 km, well below the nominal 2 km pixel size of the ABI. At this width, the contrail occupies less than half a satellite pixel, yet it produces a BTD contrast of 0.214 K at first visibility, above the 0.2 K detection threshold identified across the full SGP dataset. The observation was made on 2026-02-07 under clear-sky conditions, early morning, with no surrounding cirrus to contaminate the background. The clear sky provides a uniform background, making the narrow contrail distinguishable. The ASI image confirms that the contrail, while narrow, is optically well-defined at the time of first satellite detection. SKW4822 therefore represents the lower end of what GOES-East ABI can detect under near-optimal conditions.

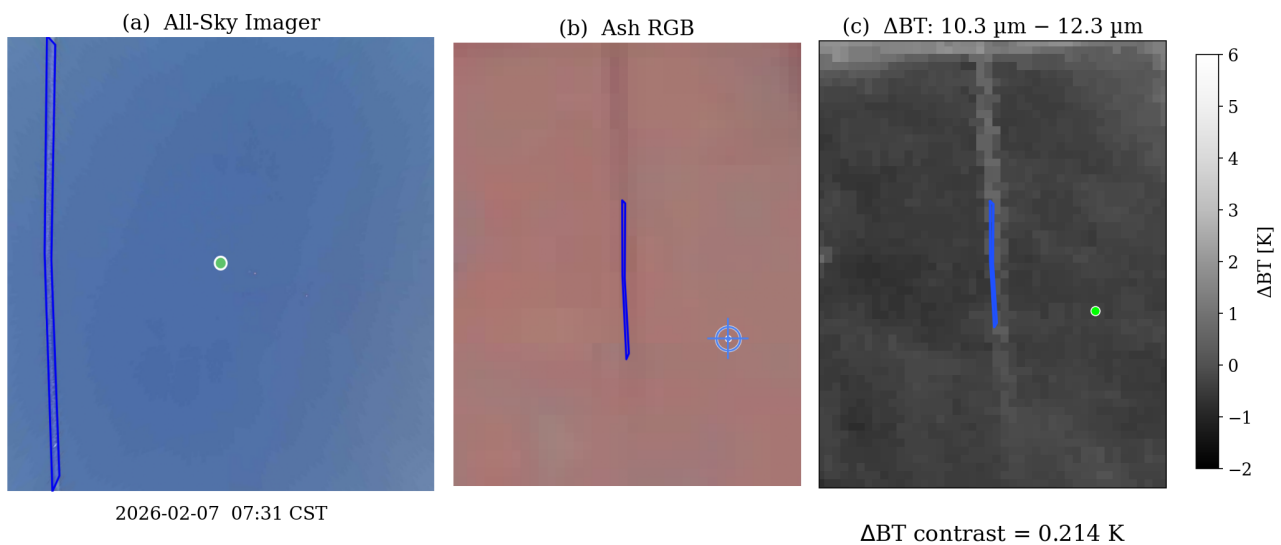


Figure 5.12: SKW4822 on 2026-02-07, maximum width 0.73 km, at the SGP site. The blue marker in the middle panel denotes the location of the camera at SGP. Despite being below the nominal 2 km pixel size, the contrail produces a positive BTD contrast under clear-sky conditions.

CAB site: MTG-I1 FCI

At the CAB site, the BTD is computed from the difference between the 10.5 μm and 12.3 μm channels at 2 km resolution of the FCI. The overall picture mirrors the SGP result: visible contrails are shifted to positive BTD contrast values and not-visible observations cluster near zero, matching the control group (Figure 5.13). However, the distributions are broader and noisier than at SGP, likely for two reasons. First, the CAB dataset is half the size of SGP, and therefore more sensitive to outliers. Second, at SGP, the Ash RGB product used for visual detection is also based on infrared channels, uses the same split-window difference technique, and has the same resolution, as the BTD analysis, so the two products are much more similar than the ones used for CAB. The Cloud Type RGB is comprised of visible and near-infrared channels, at a 1 km resolution. Roskovensky and Liou [54] show that 1.38 μm reflectance can often detect a greater amount of thin cirrus than the split-window difference BTD. Agreement between the two products is therefore not guaranteed.

Figure 5.14 shows the CAB contrail widths and BTD values at first visibility. This figure confirms the




Figure 5.13: BTD contrast distribution at CAB, separated by visibility label and observation stage, with the control group shown for comparison. The separation between visible and not-visible observations is present but weaker than at SGP.

same positive trend between contrail width and BTD contrast seen at SGP, but with systematically lower values, below 1K. This discrepancy could be attributable to the difference in resolution, but also to the different Spectral Response Function (SRF) of the two imaging radiometers. An SRF describes the response of an imaging radiometer in a particular band to radiation of different wavelengths within that band. Because the SRFs for the ABI and FCI are not identical, they can produce different BTD signatures even when observing the exact same atmospheric feature at the same bands. This means that while we expect a positive BTD signal to verify contrail presence, the magnitude of this positive signal will not be identical across BTD images by different instruments.

Four observations at CAB show a negative BTD contrast at first visibility, at widths of approximately 0.85, 1.1, 1.45, and 1.75 km. These are contrails that were labeled visible in the Cloud Type RGB, and we know they are there based on the ASI, yet produce no detectable signal in the BTD. This might be attributed to the resolution difference between the two products, and a narrow contrail may simply be too small to appear in the coarser product. To verify whether this is the case, the 10.5 μ m channel at 1 km high-resolution (HR) was downloaded. This is the only available infrared product at 1 km resolution from MTG, making it the natural choice for isolating resolution as a factor. Three case studies looking at contrails of varying widths and BTD contrast signals are analysed below, using the 3 FCI products available.

Case study: narrow contrail undetectable in BTD image — DKH1660 The contrail produced by flight DKH1660 on 10 January 2026 has a maximum width of 836m at first observability, below the 2 km resolution of the BTD channels. Looking at Figure 5.15, the contrail is visible in the left panel, the Cloud Type RGB. In the middle panel, the BTD image, the contrail is invisible. This is where the negative signal between the polygon and its surroundings was measured, and it is evident why in the image. However, the 10.5 μ m HR image resolves the contrail (right panel), confirming that the negative BTD value measured is resolution-limited: the signal exists, but the pixel resolution is too coarse to capture it.

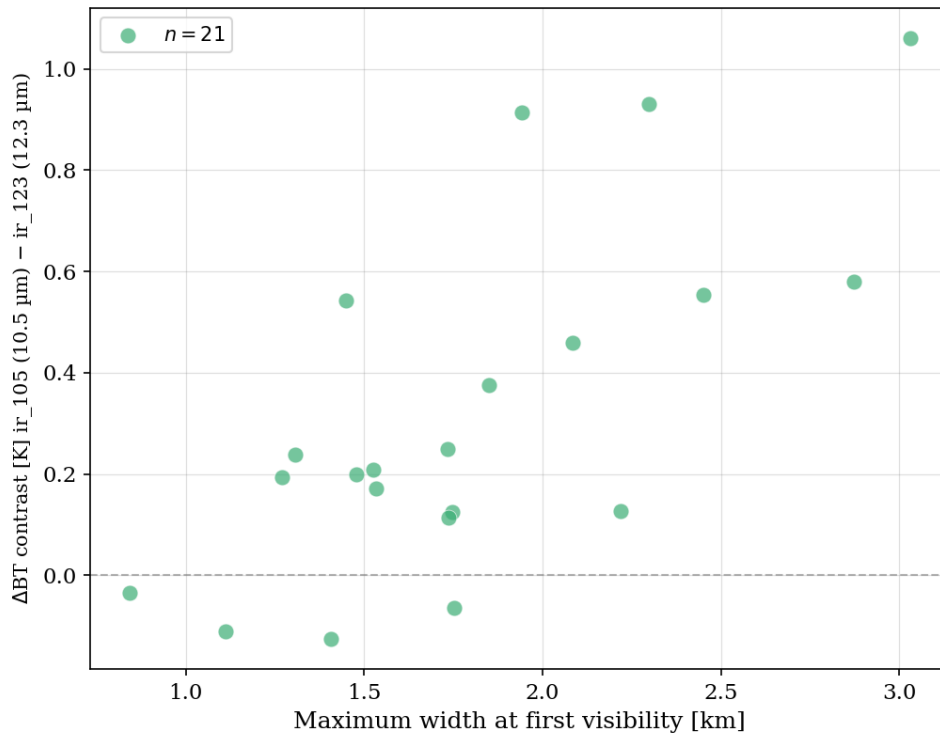


Figure 5.14: BTD contrast at first visibility of contrails at CAB.

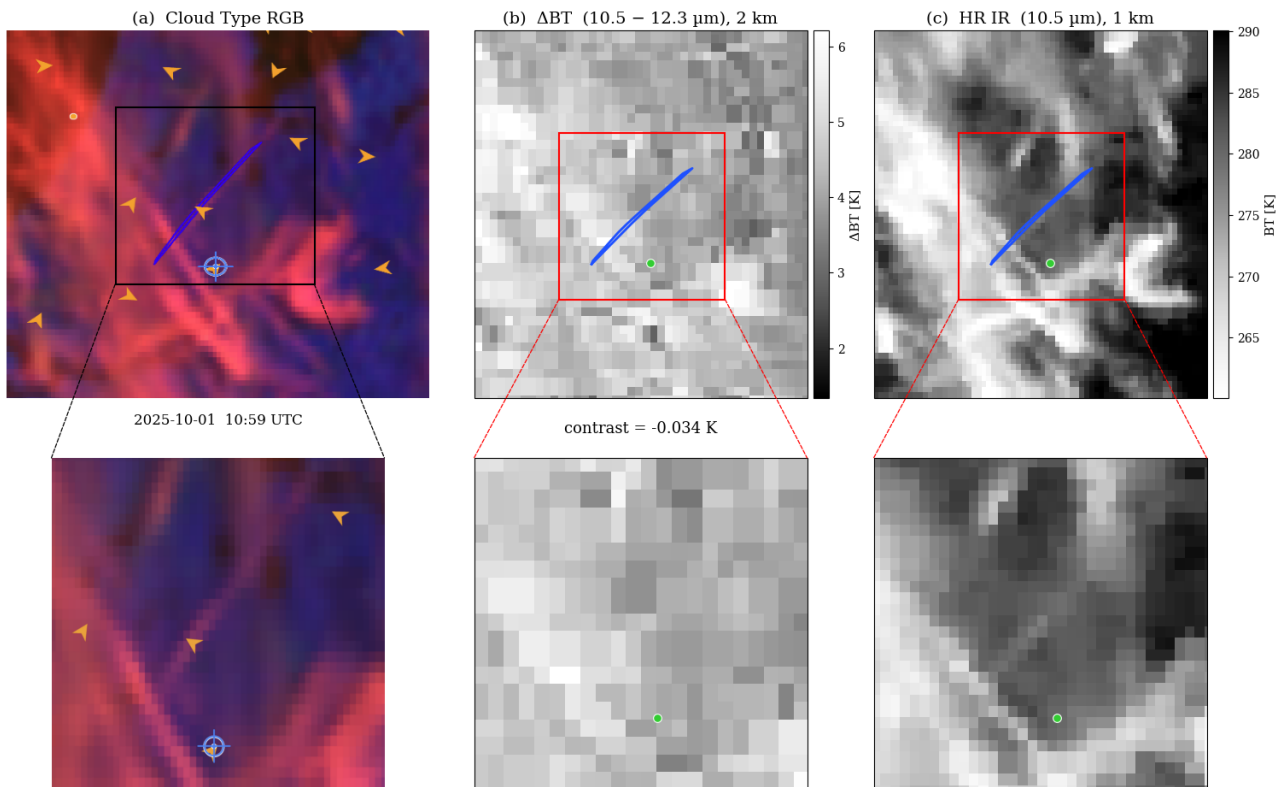


Figure 5.15: DKH1660 (width 836 m): (a) Cloud Type RGB, contrail visible; (b) BTD at 2 km, contrail absent with negative contrast; (c) HR IR 10.5 μ m at 1 km, contrail resolved.

This result stands in contrast to SKW4822 at SGP (Figure 5.12), which was detected in the ABI BTD image at a similar width of 0.73 km with a positive contrast. There are two possible explanations for this. First, once again, the larger satellite zenith angle at CAB produces a larger effective pixel

footprint, which may be too coarse to resolve a contrail of this width. Second, while the Cloud Type RGB shows the background surrounding the contrail is not directly embedded in cirrus, the sky is not as clear as in the SKW4822 case, and nearby contrails and thin cirrus may be contaminating the background BTD estimate, leading to the negative contrast value.

Case study: narrow contrail detectable in BTD image – UAL945 The contrail produced by UAL945 is one of the narrowest contrails with a positive BTD contrast of 0.238 K at first visibility in CAB (Figure 5.16). It is also one of the youngest contrails detected in the entire contrail dataset, becoming visible 8 minutes after formation. It is first detected at a maximum width of 1.305 km. In the middle panel, we can see a contrail signal in the area where we expect it, which is enough to provide a positive contrast, however without knowing the exact location of the contrail it may not have been detected by looking at the image without the ASI polygon projected upon it. It is also possible that the contrail signal is increased because of the second contrail right next to it. In the panel on the right, which represents the high resolution IR 10.5 μ m image, the contrail is once again distinguishable from its surroundings, suggesting that lower BTD contrast values seen in CAB are caused by the lower resolution of the channels used.

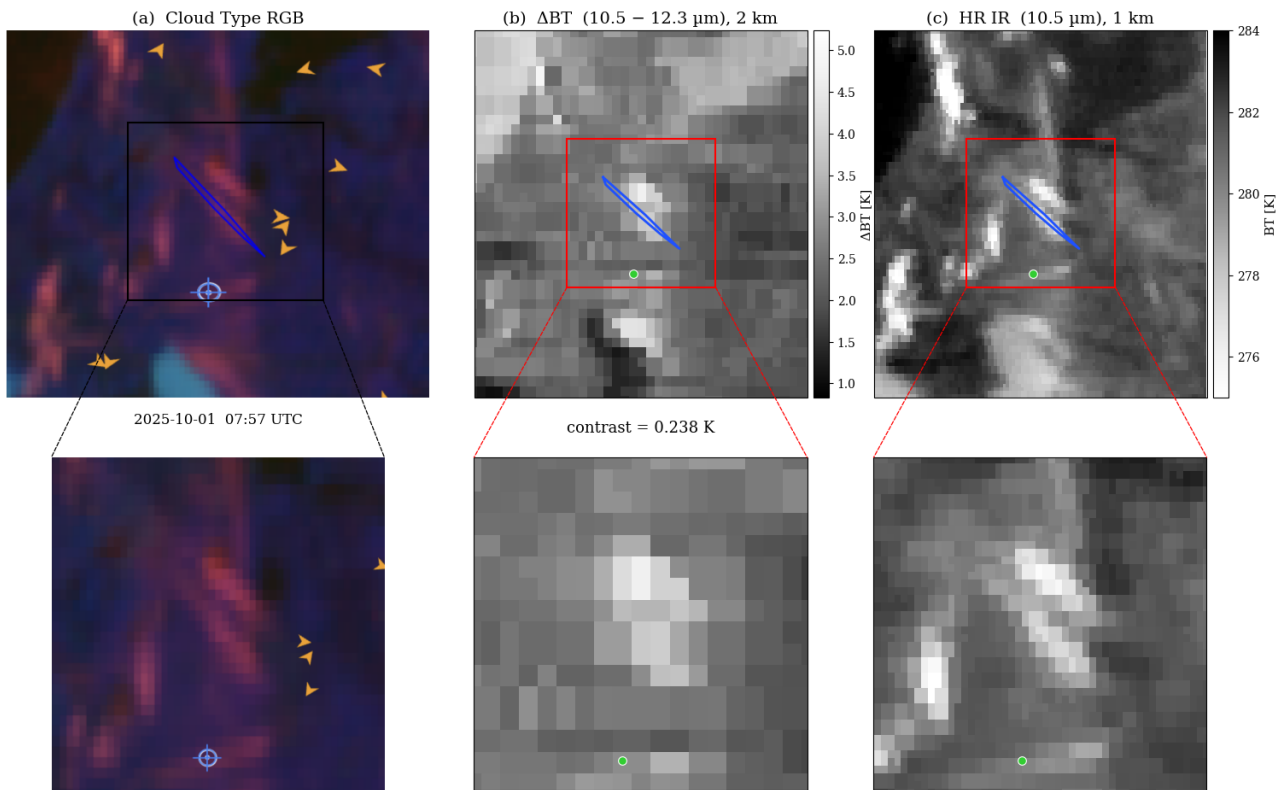


Figure 5.16: UAL945 (width 1.305 km) at first visibility: (a) Cloud Type RGB, contrail visible; (b) BTD at 2 km, positive contrast of 0.238 K despite the sub-pixel width ; (c) HR IR 10.5 μ m at 1 km, contrail resolved

Case study: young contrail undetectable in IR images – DLH453 Finally, one more case study which is interesting to look at is the contrail formed by DLH453. As can be seen in the leftmost Cloud Type RGB panel in Figure 5.17, this contrail becomes visible in a background of heavy cirrus and other contrails. At the time of first observation seen in the image, it is 1.75 km wide and 6 minutes old. It is the widest contrail with a negative BTD contrast value, most likely due to the heavy background cirrus and contrails surrounding it. As can be seen in both the middle and right panels, representing the two infrared images, the contrail is not observable, even in the

HR 10.5 μm image. This is an interesting case study because it demonstrates the Cloud Type RGB's high sensitivity to young, thin cirrus [54]. Other contrails also visible in the Cloud Type RGB are invisible to the infrared images.

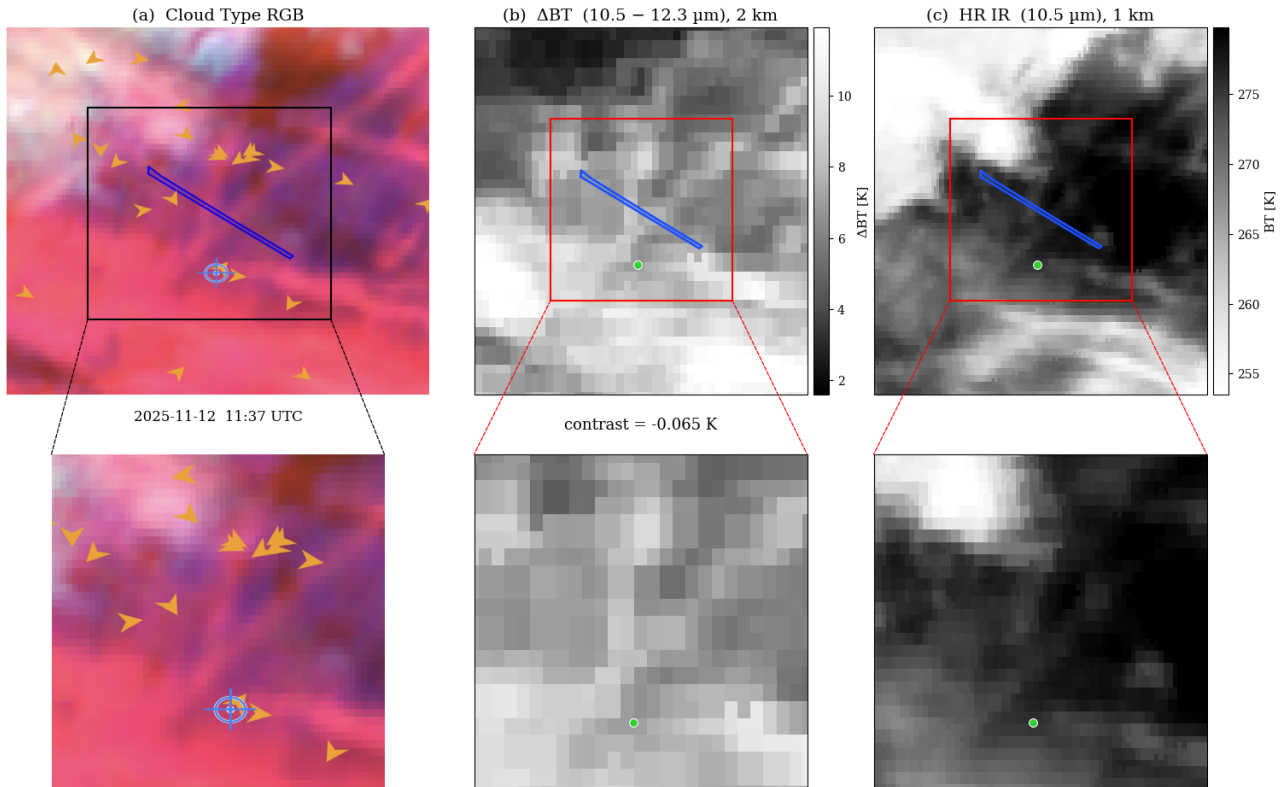


Figure 5.17: DLH453 (width 1.75 km) at first visibility: (a) Cloud Type RGB; (b) BTD image; (c) HR IR 10.5 μm image

5.3 What Affects the Delay?

The preceding analysis identified the characteristics that determine whether a contrail is detectable in a given satellite image. A contrail becomes observable once it produces a BTD contrast above the background noise level. The time it takes to reach this point defines the detection delay. Two factors that may influence how quickly this happens are analyzed: the ambient meteorological conditions at contrail altitude, which influence how rapidly the contrail spreads [55, 52] and therefore grows to a detectable width, and the background sky conditions, which determine the background brightness temperature variability the contrail signal must exceed to be detectable.

5.3.1 Age and Spreading Rate

The spreading rate \dot{w} of each contrail is computed as the mean of the instantaneous rates across all consecutive observation pairs:

$$\dot{w} = \frac{1}{N-1} \sum_{i=1}^{N-1} \frac{w_{i+1} - w_i}{t_{i+1} - t_i} \quad (5.1)$$

where w_i is the maximum width at observation i , t_i is the corresponding contrail age in minutes, and N is the total number of observations for that contrail. The result is expressed in meters per minute.

It is derived from the manual polygon annotations of contrails in the ASI and carries uncertainty. As was shown in subsection 4.1.9, the width is taken as the maximum extent of the annotated polygon, which reflects the annotator’s visual judgment of the contrail boundary, and different annotators might yield slightly different widths. Second, with some contrails having as few as two observations, a single poorly-annotated width can dominate the mean spreading rate. As a result, spreading rate estimates should be interpreted with caution: biases in annotation width will propagate directly into the computed rates, and contrails with few observations are particularly susceptible.

Across both sites, per-contrail spreading rates range from 15 to 299 m/min, with a median of 73 m/min and a mean of 88 m/min. Freudenthaler, Homburg, and Jäger [56] reported spreading rates of 18–140 m/min from ground-based lidar measurements of 14 contrails, with 65 m/min median, consistent with this study’s dataset.

Figure 5.18 shows the correlation between spreading rate and age. This is one of the primary metrics which might explain delay in observability of certain contrails in satellite images. A negative relationship is visible at both sites: contrails with higher spreading rates tend to become detectable at younger ages, while slower-spreading contrails are first detected later. This is consistent with the expectation that a faster-spreading contrail reaches a detectable width in less time, reducing the detection delay. The cases in which contrails spread very slowly yet become visible very early are ones like SKW4822 (Figure 5.12) and DKH1660 (Figure 5.15): young, narrow contrails with very high contrast against their backgrounds.




Figure 5.18: Observed contrail spreading rate versus age at first satellite detection at SGP and CAB. Each point represents one contrail.

5.3.2 Ambient Conditions

Wind Shear and Spreading Rate

Schumann et al. [57], in their compilation of over 230 observations, identifies wind shear as the primary driver of lateral contrail spreading once the wake vortex regime ends after the first ~ 100 s. Jensen et al. [55] further identify shear-driven spreading as the dominant mechanism controlling contrail width growth over time. The cross-track component of vertical wind shear was averaged across all observations of each contrail. Across both sites, cross-track shear ranges from 0.2 to 9.8 m/s/km with a mean of 3.0 m/s/km. The within-contrail standard deviation is small relative to this range (median 0.29 m/s/km), suggesting that the shear environment was approximately steady over each contrail’s observation window.

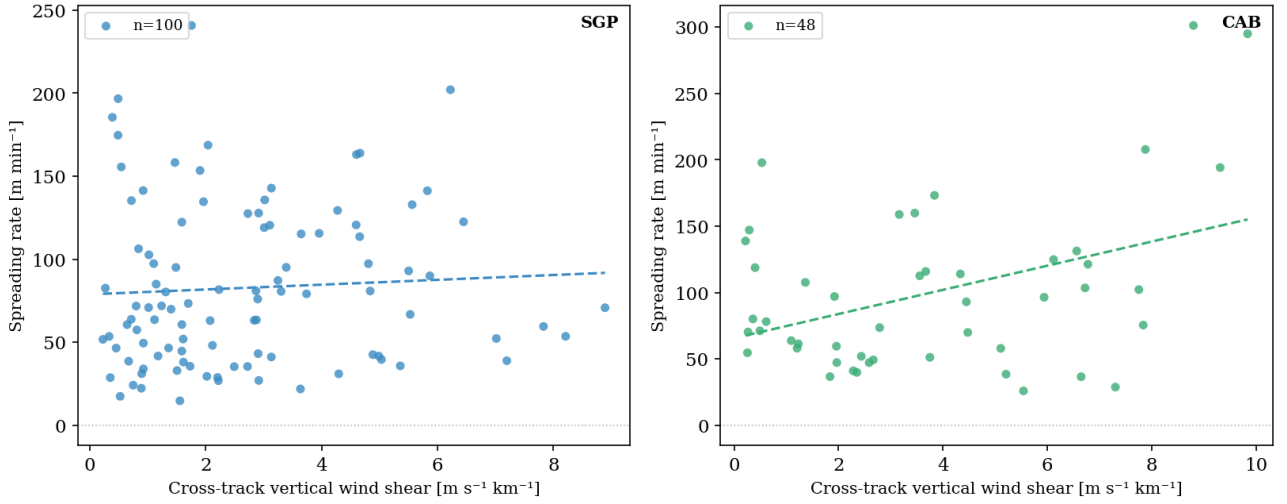


Figure 5.19: Cross-track vertical wind shear versus observed contrail spreading rate at SGP and CAB. Each point represents one contrail.

A positive trend is visible at CAB but absent at SGP, with large scatter at both sites. The difference between sites may partly reflect the disparity in sample size and number of observation days, but given the scatter present even at CAB, no significant relationship between cross-track vertical wind shear and spreading rate can be drawn from this dataset. Resolving this relationship would require higher-resolution wind data to better capture the local shear environment at each contrail, as well as direct measurements of contrail vertical and lateral extent from a ground-based lidar, which would provide a more physically grounded measure of spreading than image-derived width alone.

Atmospheric Stability and Spreading Rate

We characterize atmospheric stability at contrail altitude using the Brunt-Väisälä frequency N_{BV} , a measure of how strongly the atmosphere resists vertical air motion. In more stably stratified air, represented by higher N_{BV} , vertical displacements are suppressed, confining the contrail to a shallower layer [52]. In weakly stratified air, contrails can grow deeper vertically, exposing them to a larger range of wind speeds across different altitudes. This increased wind shear enhances lateral spreading [55]. Unterstrasser and Gierens [52] found that in strongly stable cases ($N_{BV} = 20 \times 10^{-3} \text{s}^{-1}$), vertical displacement becomes negligible. The theoretical expectation is therefore that higher N_{BV} corresponds to slower lateral spreading.

The N_{BV} values per observation in this dataset range from 3.99 to $23.7 \times 10^{-3} \text{s}^{-1}$ across both sites, consistent with the range of $5\text{--}20 \times 10^{-3} \text{s}^{-1}$ used in simulations by Unterstrasser and Gierens [52]. This range also encompasses the mean mid-latitude stability value of $N_{BV} \approx 17 \times 10^{-3} \text{s}^{-1}$ reported by Schumann et al. [57].

Schumann [10] shows that wake descent in stratified air is governed by the dimensionless product $N_{BV}t_0$, where t_0 is the aircraft-specific wake vortex timescale, determined by aircraft mass, wingspan, and speed. Heavier aircraft produce greater initial contrail depth and width, so the same ambient N_{BV} drives different scales of depth and width for different aircraft types. As the aircraft type is known for each contrail in this dataset, flights are classified by ICAO Wake Turbulence Category (WTC), sourced from the FAA Aircraft Characteristics Database [58], as either heavy (MTOW > 136 000 kg, all widebody types) or medium (narrowbody and regional jets), so that any weight-class dependence can be assessed alongside the stability relationship. The hypothesis is that heavier aircraft should produce greater lateral spreading at a given N_{BV} , as their larger wake vortices establish a deeper initial contrail layer, exposing it to a wider range of wind shear.




Figure 5.20: Mean Brunt–Väisälä frequency versus mean contrail spreading rate per contrail at SGP (left) and CAB (right), colored by ICAO wake turbulence category: dark shading denotes heavy aircraft (MTOW > 136 000 kg); light shading denotes medium aircraft. Dashed lines show linear fits to each group.

Figure 5.20 shows no clear trend between N_{BV} and spreading rate at either site, and no consistent separation between weight classes is apparent. At SGP, the heavy class trend has a negative slope, which would align with the theoretical expectation, but this group contains only 11 contrails and cannot be considered meaningful. Several factors likely obscure the underlying signal. The N_{BV} values are derived from ERA5 reanalysis data, which may not fully resolve fine-scale temperature gradients at contrail altitude. The spreading rate itself carries uncertainties arising from the polygon-based width measurement and the discrete five-minute interval between observations. Furthermore, this analysis would benefit from direct measurement of contrail vertical depth, for example from ground-based lidar, which would replace the ERA5-derived N_{BV} with the contrail depth that drives shear-induced spreading. Combined with the generally small and unbalanced sample across weight classes, these sources of uncertainty are likely masking any relationship that would otherwise be expected.

Relative Humidity over Ice

Figure 5.21 shows the RH_i distributions for contrails that became or did not become visible in geostationary satellite imagery. At CAB, both visible and never-visible observations concentrate near ice supersaturation, with visible contrails tending toward slightly higher RH_i values. At SGP, the non visible distribution peaks around $RH_i = 1.0$, but visible contrails are spread across a wider range, including drier conditions. The expectation is that visible contrails form in conditions which are more supersaturated with respect to ice, which makes them optically thicker and enables them to persist long enough to become detectable in a satellite image. The discrepancy between visible and non-visible contrails at SGP will be analyzed further in the next section. An explanation is the known unreliability of ERA5 data when predicting ice supersaturation in the context of contrail formation, as explained by Gierens, Matthes, and Rohs [59]. Despite using model-level data with 200-300 m vertical resolution at cruise altitudes, the 31 km horizontal resolution potentially smooths out localized supersaturation, and ERA5 correctly identifies ice supersaturation only 20.5 % of the time compared to in-situ measurements.

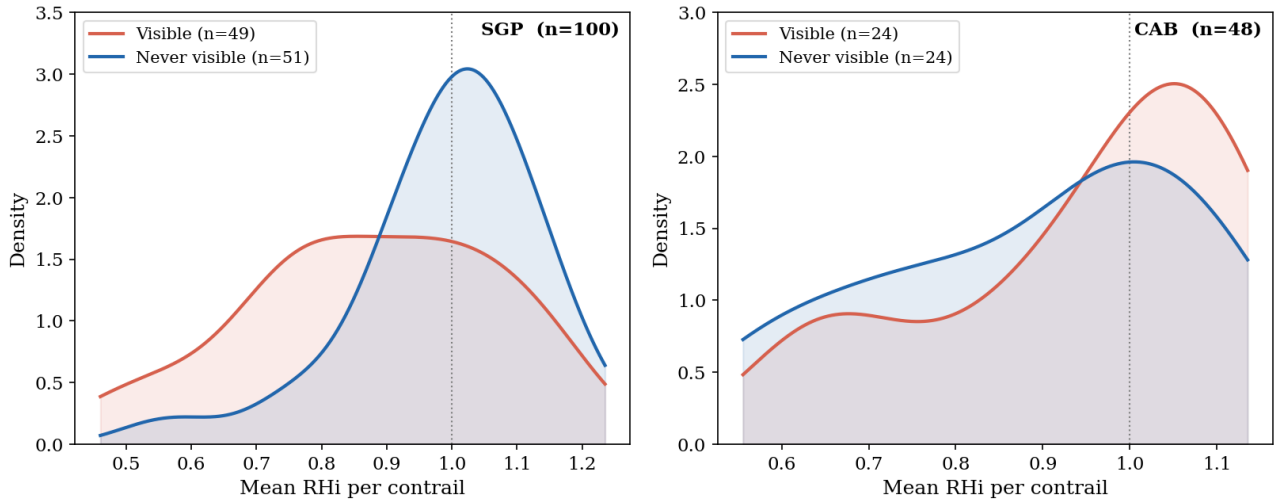


Figure 5.21: Distributions of ERA5 relative humidity over ice (RH_i) across all contrails, separated by whether the contrail became visible in satellite imagery or not, for SGP (left) and CAB (right). The dashed vertical line marks supersaturation with respect to ice ($RH_i = 1.0$).

5.3.3 Background Conditions

Table 5.4 summarizes contrail visibility rates broken down by site and sky condition. How these labels are determined is explained in subsection 4.3.4.

Site	Sky condition	Total	% Visible	% Never visible
SGP	Clear	52	55.8	44.2
	Cirrus	48	41.7	58.3
CAB	Clear	33	66.7	33.3
	Cirrus	15	20.0	80.0

Table 5.4: Contrail visibility rates by site and sky condition. Percentages are computed within each sky-condition group.

As can be seen in Table 5.4, CAB presents limitations for its analysis. There are only 6 observation days, and the split between sky conditions is uneven: four of the six days are classified as clear, producing 33 observations, while only two cirrus days contribute 15 observations. More critically, only three contrails under cirrus conditions at CAB are ever detected as visible. A subsample of three visible contrails is too small to characterize any distribution, let alone compare it to the 22 clear-sky visible contrails from the same site. The following analysis therefore considers SGP only, where both sky conditions are more evenly represented (52 clear, 48 cirrus) and the visible subgroups are large enough to support comparison.

The first property to be analyzed by sky condition is the RH_i . Looking at Figure 5.22 it can be seen that on days with cirrus, the observations are concentrated above ice supersaturation (mean $RH_i = 1.03$) while clear sky observations span a wider and drier range (mean $RH_i = 0.85$). This is consistent with literature, which confirms that natural cirrus forms mainly at higher ice supersaturation values than those necessary for contrail formation and persistence [52, 8, 43]. This figure also acts as validation of the manually selected sky condition labels. Combined with Table 5.4, this information explains the behavior seen in Figure 5.21, where not visible contrails are concentrated at higher RH_i values than visible ones. At this small sample size, it is because non-visible contrails are more common on days classified as "cirrus" (58.3% never visible), and those same days are characterized




Figure 5.22: RH_i distribution at SGP, separated by "clear" and "cirrus" sky condition labels.

by higher ambient RH_i .

Looking further at how background sky conditions might affect visibility, the left panel of Figure 5.23 shows a difference between width at first visibility in "clear" versus "cirrus" labeled observation distributions. The latter peak later, at widths of 1.5 km and above, while the former are more often visible at narrower widths. This makes sense intuitively. In the right panel of the figure, the age at first visibility is shown. Even though contrails observed against cirrus are systematically wider at first detection, they are generally detected as early in their lifetime as clear sky ones. This suggests that under cirrus, a contrail may already be present and spreading from an early age, but only becomes distinguishable from the background once it has reached a greater width. The detection delay is therefore not longer, but the width threshold that must be crossed before detection is higher.

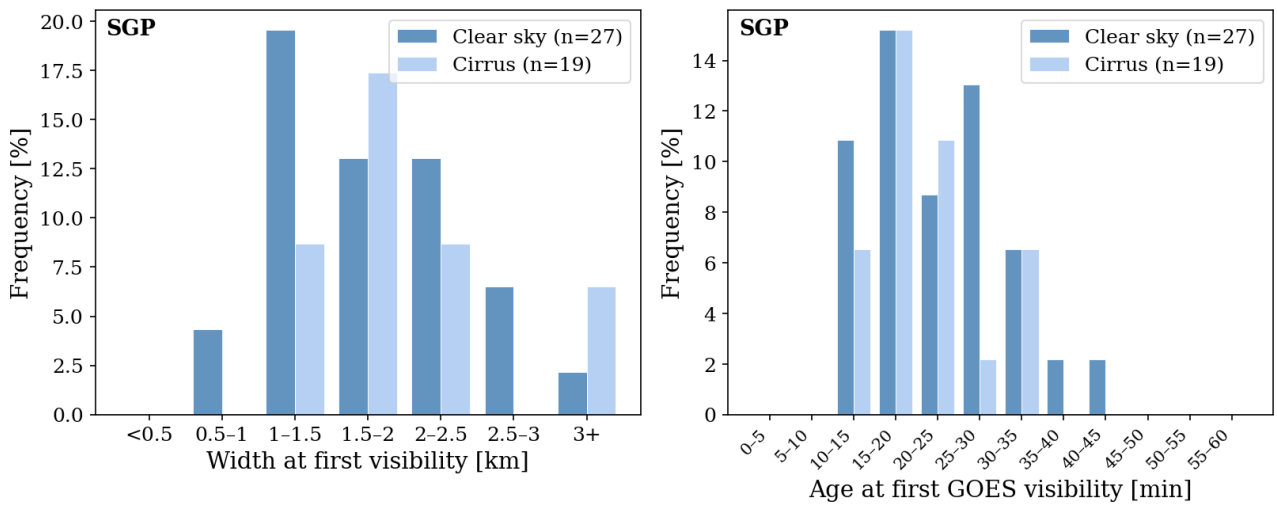


Figure 5.23: Width at first visibility for SGP (left) and age at first visibility (right), split by sky condition.

Conclusion

This thesis set out to quantify the observability of contrails in geostationary satellite imagery by constructing a dataset that bridges the gap between early contrail formation, observed from the ground, and detection from space. Two research questions guided the work: which combination of contrail properties, background conditions, and instrument characteristics determines whether a contrail appears in a given satellite image, and how long after formation a contrail takes to become observable, and what governs that delay.

To address these questions, a dataset of 148 individually tracked and flight-attributed contrails was constructed from two ground-based All-Sky Imager sites: the ARM Southern Great Plains facility in Oklahoma and the KNMI CAB Experimental Site for Atmospheric Research (CESAR) in the Netherlands. In total, 155 hours of manual observations across 40 days were collected. Using the VOCASION annotation platform, contrail polygons were drawn in rectified fisheye images at five-minute intervals, attributed to individual flights, and tracked from first appearance in the ground camera into the corresponding geostationary satellite image. For each observation, contrail width and age were estimated, ambient meteorological conditions including relative humidity with respect to ice, vertical wind shear, and atmospheric stability were interpolated from ERA5 reanalysis fields, and satellite visibility was assessed through a brightness temperature difference (BTD) contrast analysis.

Addressing the first research question, contrail width, background conditions, and instrument characteristics all contribute to whether a contrail is detectable in geostationary satellite imagery. Width at first satellite detection has a mean of 1.86 km at SGP and 1.79 km at CAB. However, the case studies show that width alone is not sufficient to determine detectability. The contrail produced by flight SKW4822 became visible at just 0.73 km under clear sky, where a clear background was enough to produce a detectable contrast despite the sub-pixel width. The contrail by JBU577 exceeded 2 km in width under clear sky but was not detectable, pointing to contrast against the background as the limiting factor rather than width. Under cirrus background conditions, contrails require greater widths to become detectable, and the visibility rate drops from 55.8% under clear sky to 41.7% under cirrus at SGP. The instrument comparison between SGP and CAB shows that the oblique viewing angle at CAB partially offsets the FCI's resolution advantage, with the DKH1660 case confirming that resolution still matters for narrow contrails: undetectable in the 2 km resolution BTD product but clearly visible in the 1 km HR IR channel.

Addressing the second research question, contrails are first detected at a mean age of 21.3 minutes at SGP and 19.0 minutes at CAB. Contrails with higher spreading rates tend to be detected at younger ages, which is consistent with the expectation that a faster-spreading contrail reaches a detectable width sooner. Contrails observed against cirrus tend to be detected at similar ages but greater widths than clear sky contrails, suggesting that a wider contrail is needed to produce a signal when the background is less uniform. No clear relationship between ERA5-derived cross-track wind shear or Brunt–Väisälä frequency and spreading rate was found, likely due to the coarse horizontal resolution of ERA5 and the limited sample size. The case studies illustrate that detection delay cannot be attributed to a single factor. The contrail produced by flight UAL945 became visible after just 8 minutes at 1.3 km width, reflecting a high spreading rate under clean background conditions. In contrast, the contrail produced by flight JBU577 was not detected until 40 minutes after formation at 2.5 km width under clear sky, where the low contrast against its background delayed detection.

6.1 Recommendations for future work

There are several directions in which this analysis could be meaningfully extended. The present dataset does not capture contrail persistence or lifetime because contrails frequently advect out of the ASI field of view before dissipating, at which point observation and annotation stops. Extending the observation window by continuing to track contrails in the satellite imagery until dissipation could yield lifetime estimates, which could then be compared with RH_i to analyze theoretical and observed persistence. Alongside this, improving the accuracy of the meteorological data would reduce uncertainty in ambient condition estimates: ERA5 operates at 31 km horizontal resolution and hourly temporal resolution, which may not resolve fine-scale variability at flight altitude [59], and in-situ measurements closer to the contrail location and altitude could produce more meaningful correlations.

Ground-based lidar could further complement this by providing direct measurements of contrail optical depth, vertical depth, and altitude. Lidar-derived altitude estimates could improve the width retrieval pipeline, which currently fixes contrail altitude from the ADS-B waypoint at first observation, and help yield more accurate spreading rate estimates. Lidar-derived contrail depth could also provide a more direct measure of vertical confinement than N_{BV} , potentially allowing a more direct test of the relationship between vertical stability and contrail spreading rate [10, 55]. Finally, direct measurement of optical depth would provide insight into an aspect of observability which this study can only discuss qualitatively.

SGP and CAB were included in this study because ASI data were available at both sites, providing a larger combined dataset. The observations at each site are taken at face value, without accounting for differences in local climate or air traffic. The two sites differ in air traffic density, with contrail coverage over Central Europe estimated at approximately 0.5% [6] compared to 0.26% over the contiguous United States [26], and in upper-tropospheric climate, with colder European flight levels reducing water vapor availability and therefore contrail optical depth [16]. These differences may contribute to the discrepancies observed between the two sites but are not further investigated here. The dataset also covers only 21 days at SGP and six days at CAB; a longer time series would improve the representativeness of the findings. Furthermore, the dataset was annotated by a single observer throughout; introducing multiple annotators could help reduce the influence of individual judgment on visibility labels and width estimates.

Finally, the dataset assembled in this study could be applied to validate contrail prediction model outputs directly, by comparing observed width evolution and age of all observations against CoCiP simulations for the same attributed flights and atmospheric conditions.

Bibliography

- [1] David S. Lee et al. “The Contribution of Global Aviation to Anthropogenic Climate Forcing for 2000 to 2018”. In: *Atmospheric Environment* 244 (2021), p. 117834. doi: 10.1016/j.atmosenv.2020.117834.
- [2] Hermann Mannstein, Peter Spichtinger, and Klaus Gierens. “A Note on How to Avoid Contrail Cirrus”. In: *Transportation Research Part D: Transport and Environment* 10.5 (2005), pp. 421–426. doi: 10.1016/j.trd.2005.04.012.
- [3] Klaus Gierens and Margarita Vázquez-Navarro. “Statistical Analysis of Contrail Lifetimes from a Satellite Perspective”. In: *Meteorologische Zeitschrift* 27.3 (2018), pp. 183–193. doi: 10.1127/metz/2018/0888.
- [4] U. Schumann et al. “Contrail Study with Ground-Based Cameras”. In: *Atmospheric Measurement Techniques* 6.12 (2013), pp. 3597–3612. doi: 10.5194/amt-6-3597-2013.
- [5] Nicolas Gourgue, Olivier Boucher, and Laurent Barthès. “A Dataset of Annotated Ground-Based Images for the Development of Contrail Detection Algorithms”. In: *Data in Brief* 59 (2025), p. 111364. doi: 10.1016/j.dib.2025.111364.
- [6] H. Mannstein, A. Brömser, and L. Bugliaro. “Ground-Based Observations for the Validation of Contrails and Cirrus Detection in Satellite Imagery”. In: *Atmospheric Measurement Techniques* 3.3 (2010), pp. 655–669. doi: 10.5194/amt-3-655-2010.
- [7] M. V. Euchenhofer et al. “Contrail observation limitations using geostationary satellites”. In: *Geophysical Research Letters* 52 (2025), e2025GL118386. doi: 10.1029/2025GL118386.
- [8] Ulrich Schumann. “On conditions for contrail formation from aircraft exhausts”. In: *Meteorologische Zeitschrift* 5.1 (1996), pp. 4–23.
- [9] Vincent R. Meijer. “Satellite-based Analysis and Forecast Evaluation of Aviation Contrails”. Doctoral dissertation. Cambridge, MA: Massachusetts Institute of Technology, June 2024.
- [10] Ulrich Schumann. “A Contrail Cirrus Prediction Model”. In: *Geoscientific Model Development* 5.3 (2012), pp. 543–580. doi: 10.5194/gmd-5-543-2012.
- [11] H. Hersbach et al. *ERA5 hourly data on pressure levels from 1979 to present*. 2023. doi: 10.24381/cds.bd0915c6. url: <https://doi.org/10.24381/cds.bd0915c6>.
- [12] Klaus Gierens, Sigrun Matthes, and Susanne Rohs. “How well can persistent contrails be predicted?” In: *Aerospace* 7.12 (2020), p. 169.
- [13] Jade Low et al. “Ground-based contrail observations: comparisons with reanalysis weather data and contrail model simulations”. In: *Atmospheric Measurement Techniques* 18.1 (2025), pp. 37–56. doi: 10.5194/amt-18-37-2025. url: <https://amt.copernicus.org/articles/18/37/2025/>.
- [14] H. Mannstein, R. Meyer, and P. Wendling. “Contrail Coverage over Western Europe Derived from NOAA-AVHRR Data”. In: *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. Vol. 4. 1999, pp. 2238–2240. doi: 10.1109/IGARSS.1999.775088.

- [15] David P. Duda et al. “A Case Study of the Development of Contrail Clusters over the Great Lakes”. In: *J. Atmos. Sci.* (2004). DOI: [https://doi.org/10.1175/1520-0469\(2004\)061<1132:ACSOTD>2.0.CO;2](https://doi.org/10.1175/1520-0469(2004)061<1132:ACSOTD>2.0.CO;2).
- [16] Ulrich Schumann. “Formation, properties and climatic effects of contrails”. In: *Comptes Rendus Physique* 6.4–5 (2005), pp. 549–565. ISSN: 1631-0705. DOI: [10.1016/j.crhy.2005.05.002](https://doi.org/10.1016/j.crhy.2005.05.002).
- [17] National Oceanic and Atmospheric Administration. *NOAA CLASS — AVHRR Data Availability*. https://www.class.noaa.gov/data_available/avhrr/index.htm. 2025.
- [18] NASA. *Moderate Resolution Imaging Spectroradiometer (MODIS)*. <https://ladsweb.modaps.eosdis.nasa.gov/missions-and-measurements/modis/>. 2001.
- [19] R. E. Murphy et al. “The Visible Infrared Imaging Radiometer Suite”. In: *Earth Science Satellite Remote Sensing*. Ed. by John J. Qu et al. Berlin, Heidelberg: Springer, 2006, pp. 199–223. DOI: [10.1007/978-3-540-37293-6_11](https://doi.org/10.1007/978-3-540-37293-6_11).
- [20] eoPortal Directory. *CALIPSO — Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations*. <https://www.eoportal.org/satellite-missions/calipso>. 2025.
- [21] National Oceanic and Atmospheric Administration (NOAA). *GOES-R Series Advanced Baseline Imager (ABI) Space Segment*. <https://www.goes-r.gov/spacesegment/abi.html>. 2025.
- [22] EUMETSAT. *MSG High Rate SEVIRI Level 1.5 Data Guide*. <https://user.eumetsat.int/resources/user-guides/msg-high-rate-seviri-level-1-5-data-guide>. 2025.
- [23] EUMETSAT. *MTG-FCI Level 1c Data Guide*. <https://user.eumetsat.int/resources/user-guides/mtg-fci-level-1c-data-guide>. 2025.
- [24] Hermann Mannstein, Richard Meyer, and Peter Wendling. “Operational detection of contrails from NOAA-AVHRR data”. In: *International Journal of Remote Sensing* 20.8 (1999), pp. 1641–1660. DOI: [10.1080/014311699212650](https://doi.org/10.1080/014311699212650).
- [25] M. Vazquez-Navarro, H. Mannstein, and B. Mayer. “An Automatic Contrail Tracking Algorithm”. In: *Atmospheric Measurement Techniques* 3.4 (2010), pp. 1089–1101. DOI: [10.5194/amt-3-1089-2010](https://doi.org/10.5194/amt-3-1089-2010).
- [26] Vincent R. Meijer et al. “Contrail Coverage over the United States Before and During the COVID-19 Pandemic”. In: *Environmental Research Letters* 17.3 (2022), p. 034039. DOI: [10.1088/1748-9326/ac26f0](https://doi.org/10.1088/1748-9326/ac26f0).
- [27] Joe Yue-Hei Ng et al. *OpenContrails: Benchmarking Contrail Detection on GOES-16 ABI*. 2023. DOI: [10.48550/arXiv.2304.02122](https://doi.org/10.48550/arXiv.2304.02122). URL: <https://doi.org/10.48550/arXiv.2304.02122>.
- [28] I. Ortiz et al. “Robust Evaluation of Neural Networks Trained on the OpenContrails Dataset”. In: *IEEE Transactions on Geoscience and Remote Sensing* 63 (2025), pp. 1–17. DOI: [10.1109/TGRS.2025.3629628](https://doi.org/10.1109/TGRS.2025.3629628).
- [29] Vincent R. Meijer et al. “Contrail Altitude Estimation Using GOES-16 ABI Data and Deep Learning”. In: *Atmospheric Measurement Techniques* 17.20 (2024), pp. 6145–6162. DOI: [10.5194/amt-17-6145-2024](https://doi.org/10.5194/amt-17-6145-2024).
- [30] Gabriel Jarry et al. “GVCCS: A Dataset for Contrail Identification and Tracking on Visible Whole Sky Camera Sequences”. In: *Earth System Science Data Discussions* (2025), pp. 1–30. DOI: [10.5194/essd-2025-444](https://doi.org/10.5194/essd-2025-444).

- [31] Rémi Chevallier et al. “Linear Contrails Detection, Tracking and Matching with Aircraft Using Geostationary Satellite and Air Traffic Data”. In: *Aerospace* 10.7 (2023), p. 578. DOI: 10.3390/aerospace10070578.
- [32] Scott Geraedts et al. *A Scalable System to Measure Contrail Formation on a Per-Flight Basis*. arXiv:2308.02707. 2023. DOI: 10.48550/arXiv.2308.02707. URL: <https://doi.org/10.48550/arXiv.2308.02707>.
- [33] Aaron Sarna et al. “Benchmarking and Improving Algorithms for Attributing Satellite-Observed Contrails to Flights”. In: *Atmospheric Measurement Techniques* 18.14 (2025), pp. 3495–3532. DOI: 10.5194/amt-18-3495-2025.
- [34] Edward Gryspeerdt et al. “Operational Differences Lead to Longer Lifetimes of Satellite Detectable Contrails from More Fuel Efficient Aircraft”. In: *Environmental Research Letters* 19.8 (2024), p. 084059. DOI: 10.1088/1748-9326/ad5b78.
- [35] Oliver G. A. Driver, Marc E. J. Stettler, and Edward Gryspeerdt. “Factors Limiting Contrail Detection in Satellite Imagery”. In: *Atmospheric Measurement Techniques* 18.5 (2025), pp. 1115–1134. DOI: 10.5194/amt-18-1115-2025.
- [36] Atmospheric Radiation Measurement (ARM) User Facility. *SGP Central Facility, Lamont, OK (C1)*. <https://armgov.svcs.arm.gov/capabilities/observatories/sgp/locations/C1>. 2025.
- [37] KNMI. *CESAR Observatorium*. <https://www.knmi.nl/kennis-en-datacentrum/project/cesar-observatorium>. 2025.
- [38] Google Research. *Contrails Explorer*. <https://contrails.webapps.google.com/main>. 2026.
- [39] T. J. Schmit et al. “A Closer Look at the ABI on the GOES-R Series”. In: *Bulletin of the American Meteorological Society* 98 (2017), pp. 681–698. DOI: 10.1175/BAMS-D-15-00230.1.
- [40] NOAA Office of Satellite and Product Operations. *GOES Schedules and Scan Sectors*. <https://www.ospo.noaa.gov/operations/goes/schedules.html>.
- [41] Luke Kulik. “Satellite-based detection of contrails using deep learning”. MA thesis. Massachusetts Institute of Technology, 2019.
- [42] Hermann Mannstein, Richard Meyer, and Peter Wendling. “Operational detection of contrails from NOAA-AVHRR-data”. In: *International Journal of Remote Sensing* 20.8 (Jan. 1999), pp. 1641–1660. DOI: 10.1080/014311699212650.
- [43] Patrick Minnis et al. “Contrail Frequency over the United States from Surface Observations”. In: *Journal of Climate* 16.21 (2003), pp. 3447–3462. DOI: 10.1175/1520-0442(2003)016<3447:CFOTUS>2.0.CO;2.
- [44] EUMETSAT. *MTGFCI Level 1C Data Guide*. <https://user.eumetsat.int/resources/user-guides/mtg-fci-level-1c-data-guide>.
- [45] *Cloud Type RGB Quick Guide*. Tech. rep. EUMeTrain / EUMETSAT, 2021. URL: <https://eumetrain.org/sites/default/files/2021-05/CloudTypeRGB.pdf>.
- [46] Hersbach, H. and Bell, B. and Berrisford, P. and et al. *ERA5 hourly data on model levels from 1940 to present*. 2023. DOI: 10.24381/cds.adbb2d47. URL: <https://cds.climate.copernicus.eu/datasets/reanalysis-era5-complete>.
- [47] Juho Kannala and Sami S. Brandt. “A Generic Camera Model and Calibration Method for Conventional, Wide-Angle, and Fish-Eye Lenses”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28.8 (2006), pp. 1335–1340. DOI: 10.1109/TPAMI.2006.153.

- [48] OpenCV. *Fisheye camera model*. https://docs.opencv.org/3.4/db/d58/group__calib3d__fisheye.html. 2024.
- [49] Sean Gillies et al. *Shapely: manipulation and analysis of geometric objects*. 2024. DOI: 10.5281/zenodo.5597138. URL: <https://doi.org/10.5281/zenodo.5597138>.
- [50] D. Sonntag. “Advancements in the field of hygrometry”. In: *Meteorologische Zeitschrift* 3.2 (1994), pp. 51–66.
- [51] Marc Shapiro et al. *pycontrails: Python library for modeling aviation climate impacts*. Version v0.61.0. 2026. DOI: 10.5281/zenodo.19597638. URL: <https://doi.org/10.5281/zenodo.19597638>.
- [52] Simon Unterstrasser and Klaus Gierens. “Numerical simulations of contrail-to-cirrus transition – Part 1: An extensive parametric study”. In: *Atmospheric Chemistry and Physics* 10.4 (2010), pp. 2017–2036. DOI: 10.5194/acp-10-2017-2010.
- [53] R. Chevallier et al. “Linear Contrails Detection, Tracking and Matching with Aircraft Using Geostationary Satellite and Air Traffic Data”. In: *Aerospace* 10.7 (2023). DOI: <https://doi.org/10.3390/aerospace10070578>.
- [54] J. K. Roskovensky and K. N. Liou. “Detection of thin cirrus using a combination of 1.38- μm reflectance and window brightness temperature difference”. In: *Journal of Geophysical Research: Atmospheres* 108.D18 (2003), p. 4570. DOI: 10.1029/2002JD003346.
- [55] Eric J. Jensen et al. “Spreading and growth of contrails in a sheared environment”. In: *Journal of Geophysical Research: Atmospheres* 103.D24 (1998), pp. 31557–31567. DOI: 10.1029/98JD02594.
- [56] V. Freudenthaler, F. Homburg, and H. Jäger. “Contrail observations by ground-based scanning lidar: Cross-sectional growth”. In: *Geophysical Research Letters* 22.24 (1995), pp. 3501–3504. DOI: 10.1029/95GL03549.
- [57] Ulrich Schumann et al. “Properties of individual contrails: a compilation of observations and some comparisons”. In: *Atmospheric Chemistry and Physics* 17.1 (2017), pp. 403–438. DOI: 10.5194/acp-17-403-2017.
- [58] Federal Aviation Administration. *Aircraft Characteristics Database*. 2024. URL: https://www.faa.gov/airports/engineering/aircraft_char_database.
- [59] Klaus Gierens, Sigrun Matthes, and Susanne Rohs. “How Well Can Persistent Contrails Be Predicted?” In: *Aerospace* 7.12 (2020). DOI: 10.3390/aerospace7120169.