Contrail Observation on Geostationary Satellite Data

ECMWF-ESA Machine Learning Workshop

Ana-Maria Pelin
Bologna, Italy
April 2026

EUROCONTROL

Agenda

1. Role of Contrail Observation
2. Observation Sources
3. Research Trends in Deep Learning
4. Instance Segmentation
5. Models Tested
6. Performance
7. Model Distillation
8. Results

What are contrails?

- Contrails or condensation trails form from the exhaust of aircraft travelling at cruise level.
- Persistent contrails form in ice supersaturated regions (ISSR) and they trap outgoing longwave radiation and greenhouse gases (GHG).

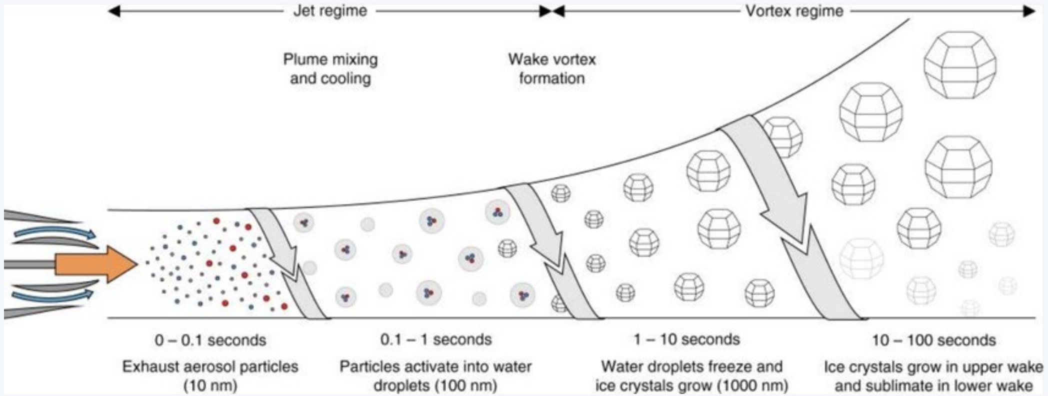

- Short-lived but globally significant in terms of aviation emissions, around 30% of CO2 climate impact.
- Two mitigation pathways: in-flight contrail avoidance and novel engine development.
- This work focuses on novel deep learning and computer vision methods for detection, modelling and tracking at observation level.

Observation Sources

Geostationary Observations

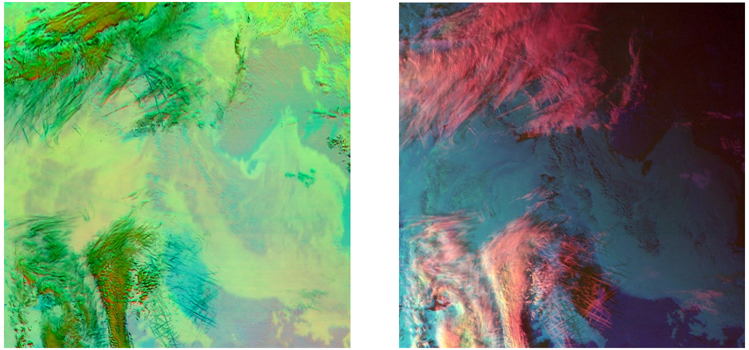

Source: Eurocontrol 2026

The Ash and Cloud Composites: AshRGB widely adopted for geostationary monitoring on GOES-16 and MTG satellites. Cloud composite works on visible short IR wavelength and possesses higher resolution.

Ground Camera Observations

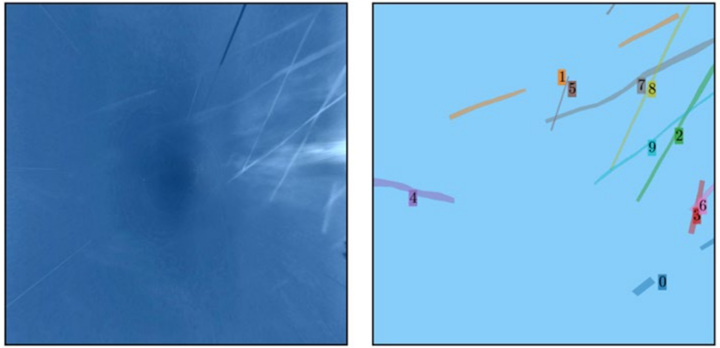

Source :Eurocontrol 2026

Ground Cameras: Local sensors (visible and IR spectrum) used to detect contrail source, persistance and dispersion mechanism. Available at EUROCONTROL Innovation Hub.

SOTA in Contrail Detection and Tracking

FOUNDING MEMBER sesar JOINT UNDERTAKING CICONIA EUROCONTROL

Geostationary Observations

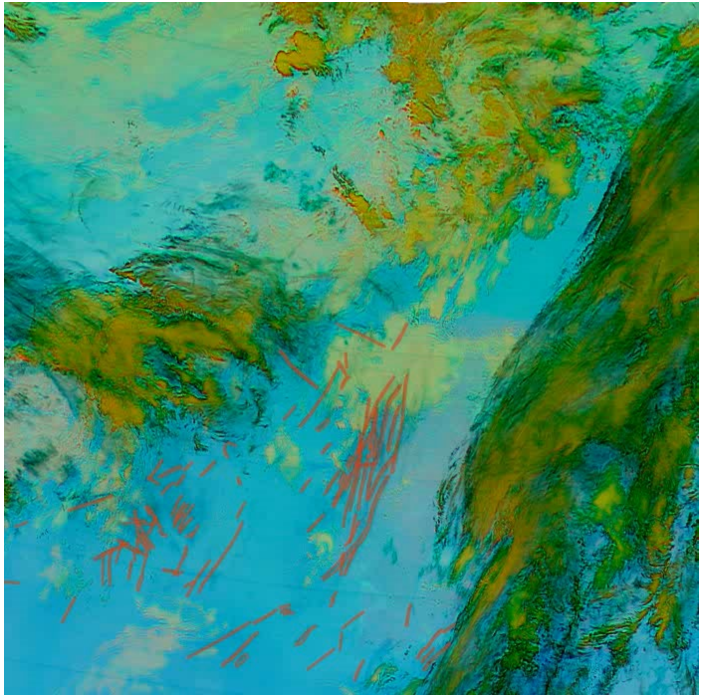

Ground Camera Observations

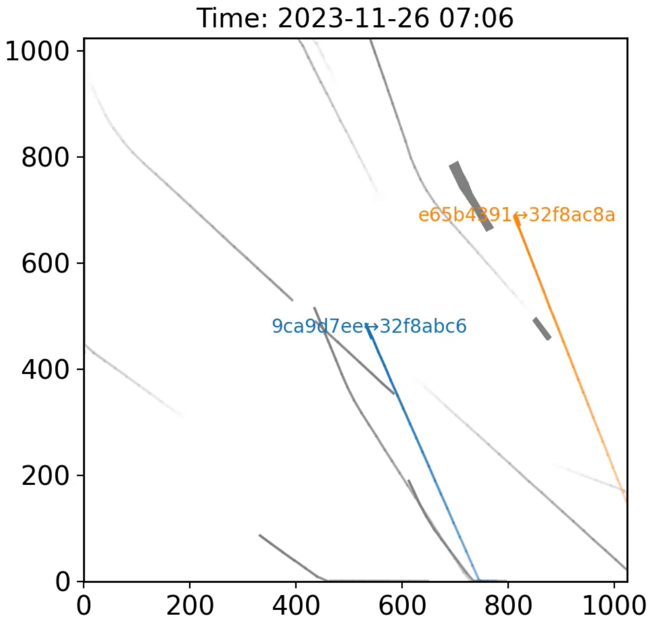

Source :Eurocontrol 2025

Newest DL-based Detection Framework

📷 Ground Camera Observations

Best Solution: Fine-tuning/adapting the SAM2 model specifically for video segmentation and tracking of contrails using ADS-B data **as object prompts**.

The Result: Segmentation, attribution and tracking are performed all at once by a single algorithm. Easier attribution as physics-informed prompt naturally solves the complex problem.

Translating to Satellite Data: Still poses large challenges

- Provide relevant object prompts,
- Parallax and overlap of flights due to low resolution,

Ideally - use the ground detections as prompts.

Geostationary and Ground Camera Observations

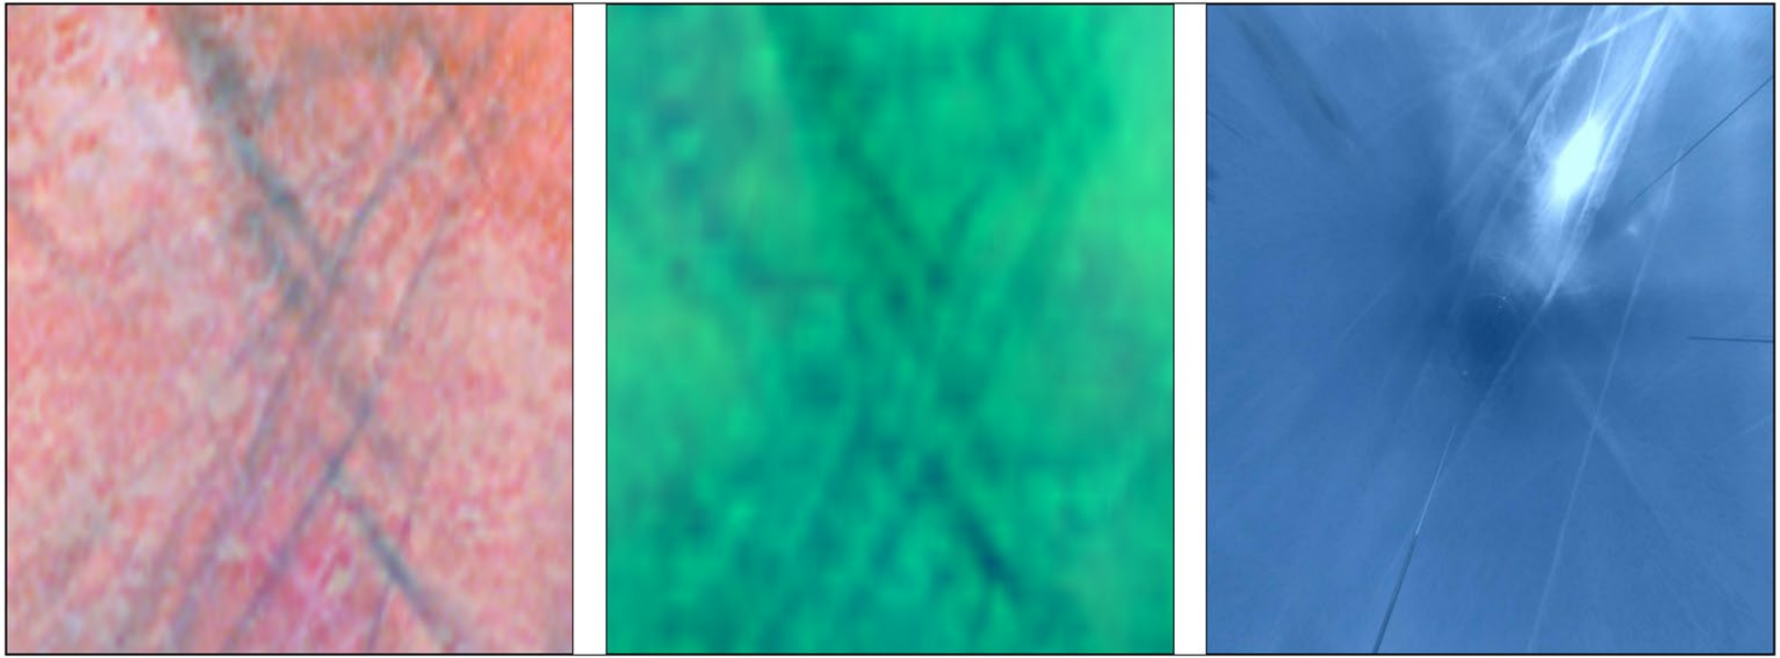

Source : Eurocontrol 2025

Instance Segmentation

Geostationary Observations

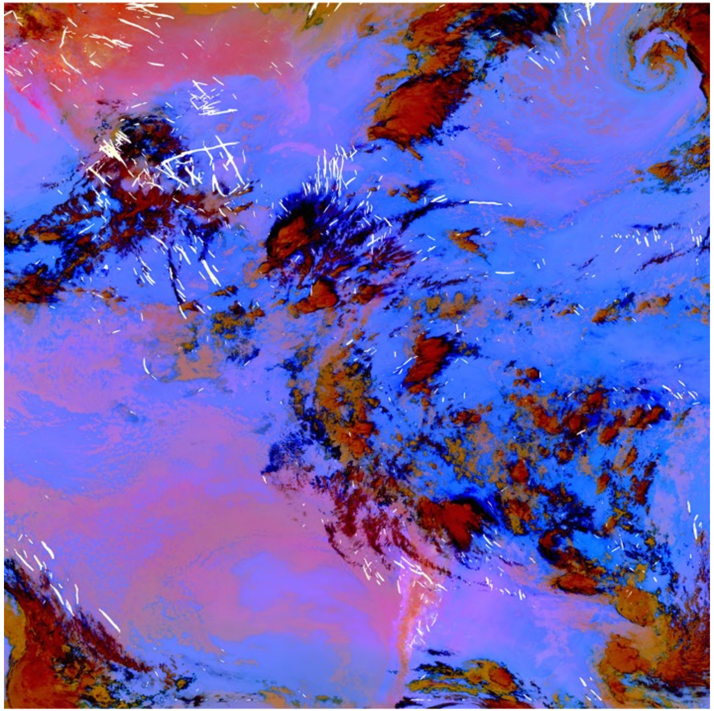

Source : Ortiz et al. et al., 2025

Best Reference Solution: Solving segmentation at a mask level with semantic segmentation models and splitting the contrails with skeletonization.

Alternative: Obtain directly the instance level masks to have a contrail count and direct ID references for easier tracking.

Goal: Understand applicability and performance of **off-the-shelf detection/segmentation** models on satellite data, to provide better masks for future attribution.

Dataset Construction

Semantic Labels (Original)

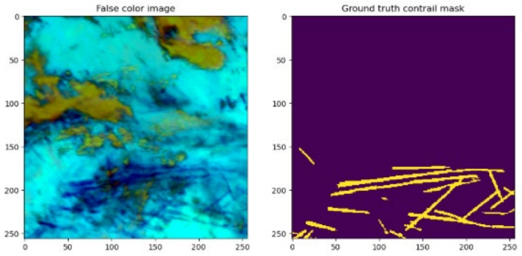

Binary masks from Kaggle GOES-16 dataset; 1 mask per sample

Source : Kaggle, Google, 2024

Instance Labels (Original)

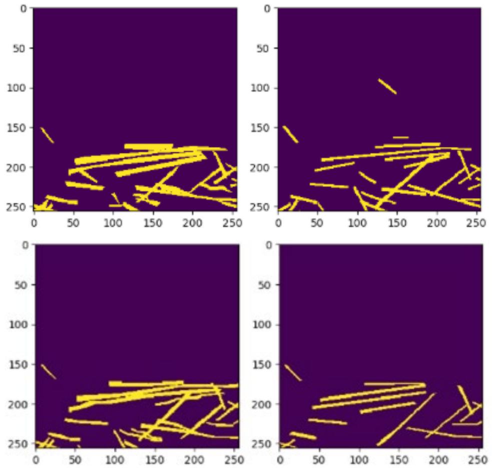

Multiple annotators, variations that are accounted by 4 individual sets of labels

Instance Detection Methods

Overview of Performance

END-TO-END DETECTORS

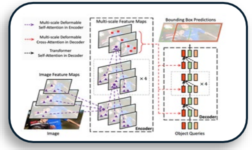

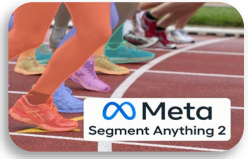

BASELINE ALGORITHM BASED ON U-NET AND MANNSTEIN

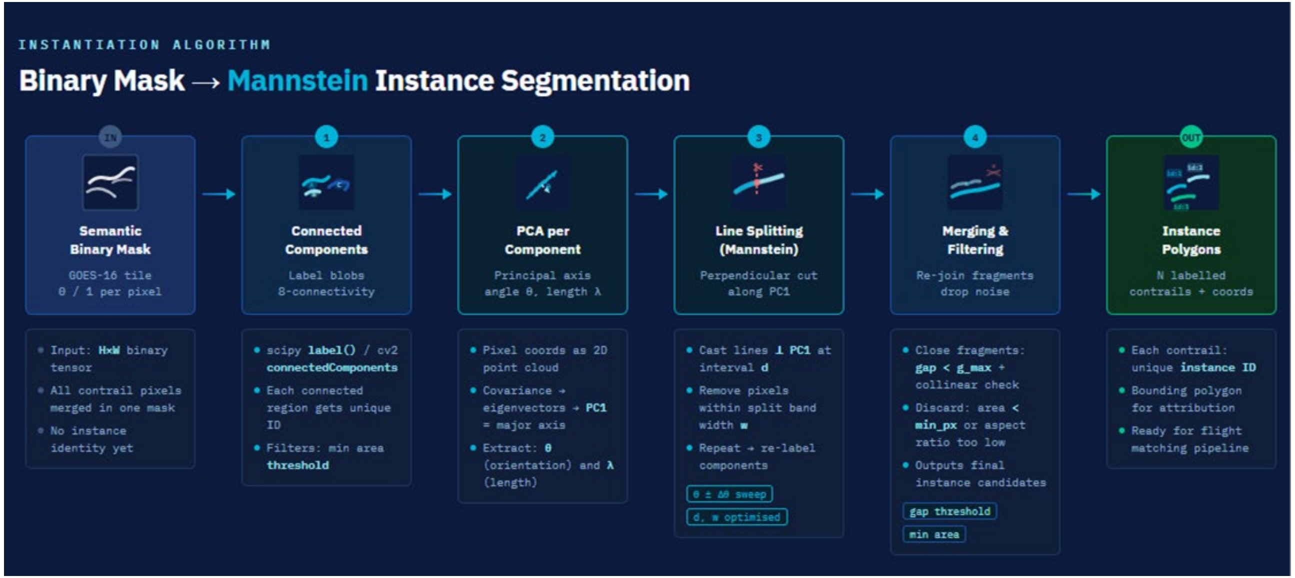

Source :Eurocontrol 2026

Models Evaluated

Architectural Features

YOLO

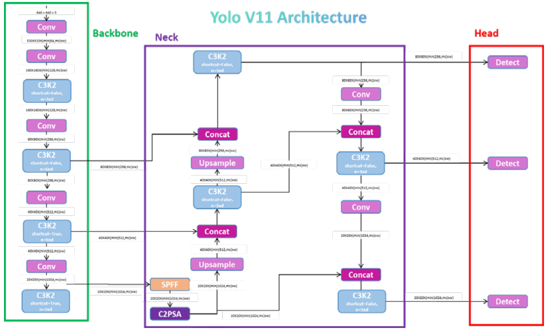

Source: Ultralytics

RF-DETR

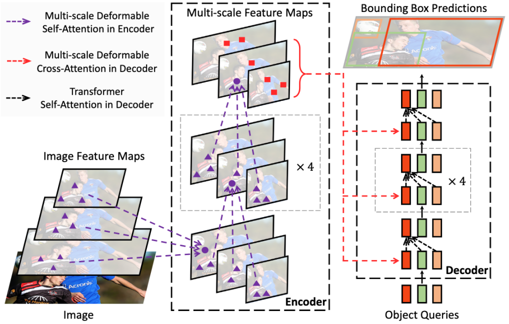

Source: Roboflow

Mask2Former

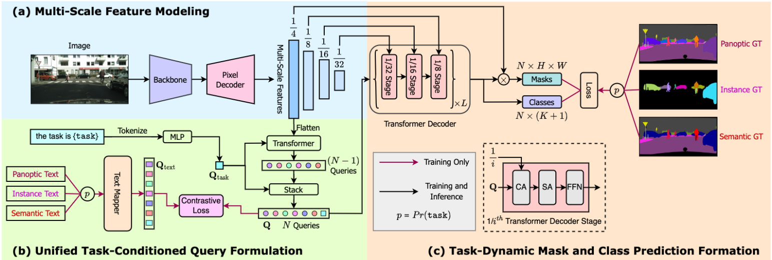

Source: Meta

UNet/UNet++

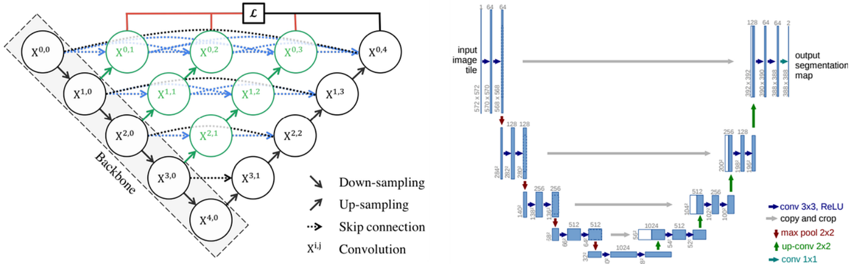

Source: Medium

Knowledge Distillation from DINOv3

Leveraging Satellite-pretrained Latent Spaces

Distillation Architecture

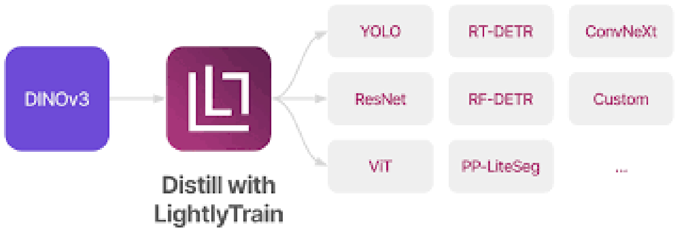

Source : LightlyAI, 2025

How it works

Teacher: Frozen DINO ViT pretrained on satellite imagery produces rich feature representations

Student: YOLO backbone trained to mimic teacher embeddings via Lightly SSL

Benefit: Better generalisation on limited contrail-labelled data

Visual Comparison

Undistilled Training of YOLO-seg & YOLO-obb

Simple training under default config.

Training of YOLO-Seg & YOLO-obb Pretrained from DINOv3 base model

Shows increase in segmentation counts. Boost not that prominent for obb models.

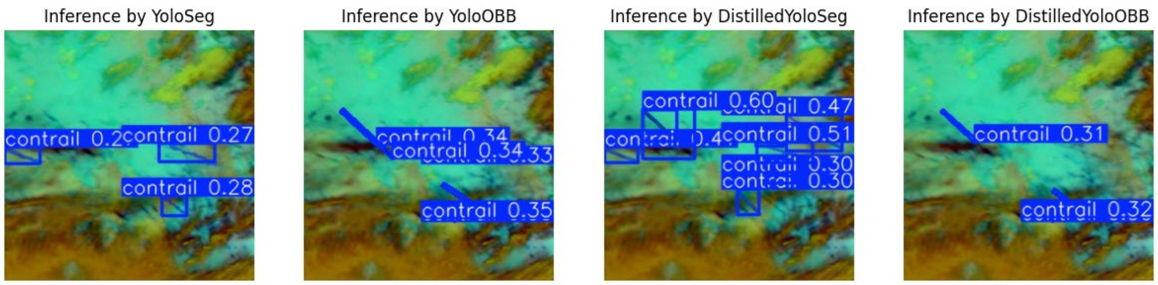

Source : Eurocontrol, 2026

Comparative Results

Instance Segmentation Performance

SESAR Joint Undertaking Founding Member · CICONIA · EUROCONTROL

Model	Dice Score	Avg IoU	mAP@50	mAP (avg)	mAP@75
UNet (Semantic)	0.6687	—	—	—	—
Mannstein (Instance)	—	0.2839	0.3507	0.3197	0.1241
YOLO Sat	**0.4754**	0.1845	0.2040	0.1938	0.0246
YOLO Distilled	**0.4636**	0.1731	0.1965	0.1791	0.0146
YOLO SegDistilled	0.4452	0.1675	0.1897	0.1759	0.0194
YOLO DistillSegLong	0.3709	0.1643	0.2007	0.1767	0.0185
YOLO Distilled Satellite	**0.4958**	0.1968	**0.2369**	0.2109	0.0141
YOLO OBB	0.0834	0.1653	0.2161	0.1853	0.0090
Yolo Seg	0.174	0.1578	0.1576	0.1523	0.0000

Key Findings

- **YOLO SAT** achieves best overall mAP of 0.237 and Dice Score of 0.4958, improving the performance of the undistilled models by 30% and the performance of the models distilled on regular images by 7%.
- Larger models show diminishing returns on thin contrail structures.
- Baseline still wins with a Dice Score of 0.6687 and a mAP@0.5 score of 0.35.

Conclusions

Key Takeaways

- Accurate detection at instance level remains challenging on satellites due to the late appearance of contrails;
- Knowledge distillation may be useful as a means of improving detection and tracking perfromance of pre-trained models on EO data, but future work should tackle it at scale;
- Future work should focus on finding reliable prompts or mapping mechanisms for colocation.

Geostationary and Ground Camera Observations

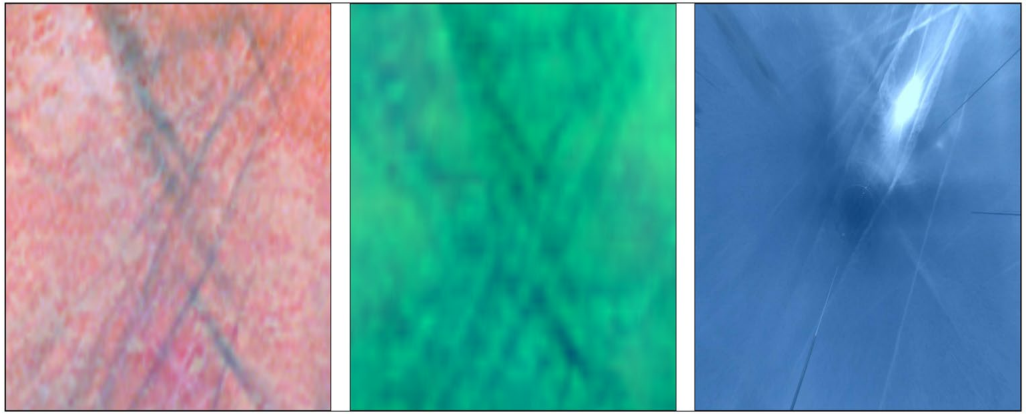

Source : Eurocontrol 2025

Thank you for your attention!

ECMWF/ESA Machine Learning Workshop

a.m.pelin@tudelft.nl
philippe.very@eurocontrol.int
ramon.dalmau-codina@eurocontrol.int
gabriel.jarry@eurocontrol.int

Ana-Maria Pelin
Bologna, Italy
April 2026