Response to Reviewer 1's Comments

Green: Reviewer's comments

Black: Authors' response

Blue: Edits in manuscript

The Authors train a machine learned (ML) model on a classification task (contrail absent, short-lived contrail, persistent contrail) using as model input imagery from aircraft-mounted rear-pointing cameras from two recent contrail study campaigns that involved a chase plane studying contrail properties and measuring ambient weather conditions. The labels for the classification task training come from applying Schmidt-Appleman criteria and RHi > 100% to the chase plane's ambient atmospheric conditions measurements. The Authors achieve F1 scores between 0.83 and 0.93 when evaluating this model in a k-fold heldout cross validation scheme.

In its current state I cannot recommend accepting the manuscript; I am on the fence and defer to the editors about whether subsequent major revisions should have a chance to be reconsidered..

We would like to thank the reviewer for their very constructive feedback. The responses to reviewer's suggestions and the associated revisions in the manuscript are provided below. We have revised the Introduction to emphasize the value and limitations for onboard camera observations of very early contrail development. Our intention was to make the case for the usefulness of these observations and to discuss considerations for future camera installations in measurement campaigns and potentially on selected commercial flights and within this context, machine learning provides an example of analyzing the imagery.

We have revised the Introduction to emphasize the value and limitations of onboard camera observations of very early contrail development. We have also revised the classification evaluation to group images by flight day and added a comparison of SAC predictions with manually assessed contrail occurrence in ND-MAX imagery, taking the camera observations as the reference.

We have also revised the classification evaluation to group images by flight day and added a comparison of SAC predictions to manually assessed contrail occurrence in ND-MAX imagery. Our responses to the specific comments and the associated manuscript revisions are provided below.

---

1. The case for scientific significance is not compellingly made:

While the manuscript is often clear enough and reproducible, the scientific significance of these F1 scores is not discussed in a compelling way. The manuscript mentions possible use of aircraft-mounted cameras as possibly decreasing expense and therefore accelerating rollouts and availability of data relevant to contrail weather forecasting, but the dots are not clearly connected: contrail forecasting is an important area to improve, but no evidence is provided (nor citations are made) for the assertion that rear-pointed camera data will improve the accuracy of contrail forecasts.

We agree that the original manuscript did not adequately distinguish the motivation for additional observations from evidence of improved forecast accuracy. The F1 scores measure the agreement between the classifier's predictions and the reference labels that were calculated from the chase aircraft's atmospheric measurements using the SAC and a threshold for relative humidity with respect to ice. They do not demonstrate improved contrail forecasts or predict the subsequent lifetime of the observed contrails. We have therefore revised the framing to emphasize the observational contribution of the work. The two campaigns allow onboard imagery of very early contrail development to be examined alongside atmospheric measurements. The comparison documents the filtering and synchronization to use these observations, the factors that limit contrail visibility, and considerations for future camera installations, either in future measurement campaigns or ultimately for some commercial flights. Image classification is presented within this context as an example of analyzing the imagery.

We have also added a comparison of SAC predictions with manually assessed contrail occurrence in ND-MAX imagery, taking the camera observations as the reference. This comparison examines whether formation predicted from the measured atmospheric conditions agrees with visible contrail occurrence (assessed manually) and it provides an observational assessment separate from the classifier's agreement with the labels but it does not test whether incorporating the observations improves weather forecast.

The revised Introduction also included references that strengthen the motivation for additional observations through discussion of Brauer et al. (2021), Schumann et al.(2017) and Low et al.(2025). These studies provide context to examine young contrail properties and compare observations with model predictions. We do not present them as evidence that our camera system improves forecast accuracy.

Section Introduction:

1. We added the following references: Brauer et al. (2021) on airborne measurements of young contrail properties and expanded the discussions of Schumann et al. (2017) and Low et al. (2025) on observational datasets, comparisons with model predictions and limitations affecting camera observations.

   Airborne measurement campaigns complement these observations by allowing contrails to be examined with in situ atmospheric measurements and provide information about their early development and the conditions in which they form.

   Brauer et al. (2021) investigated the dependence of contrail ice properties and extinction on ambient temperature and humidity during the ECLIF II/ND-MAX campaign. They found incomplete soot activation into ice particles at lower altitudes and at temperatures close to the Schmidt–-Appleman formation threshold. Their measurements showed how temperature, RHi and water vapor concentration influence ice number concentration and particle size and thus the extinction of young contrails.

   Schumann et al. (2017) compiled airborne, satellite and ground-based measurements of individual contrails with known ages, which they compared to results from CoCiP. Their database included microphysical and geometric contrail properties using in situ past contrail measurements with optical particle spectrometers, visual observations and remote sensing data from satellites, cameras as well as ground-based and airborne lidar observations, for individual aircraft contrails. In their work the authors also emphasized that testing of models requires more reliable data on contrail properties and age, atmospheric conditions, the generating aircraft, observational methods and in general additional measurements and more complete datasets.

   Cameras onboard aircraft offer an opportunity to observe the initial stages of contrail formation. When interpreted alongside in situ atmospheric measurements, these observations can be used to examine the relationship between visible contrail formation and what we expect from thermodynamics (National Academies of Sciences, Engineering, and Medicine, 2025). Their usefulness also depends on the available view of the aircraft wake, image quality and the processing to account for factors that affect contrail visibility, which in turn make camera placement and orientation important considerations for future measurement campaigns.

2. We revised “Current scope and contributions”:

   We investigate whether a machine learning model can distinguish thermodynamically defined categories from the imagery and also the observational limitations encountered

when analyzing contrails during their very early development. The main contributions of this work are as follows:

- We describe the processing and filtering of onboard camera imagery from the two NASA campaigns and document visual obstructions, limitations in image quality and the available view of the aircraft wake early in the contrail development.
- We investigate the use of image classification to detect contrail formation and identify potential ISSRs. We evaluate classification performance against reference labels from the DC-8's atmospheric measurements using the SAC and a threshold for RHi.
- We discuss limitations of the camera observations and considerations for camera mounting position and orientation for future campaigns, as the campaigns we consider in this paper differ in aircraft, engines, atmospheric conditions, and image acquisition

The longer-term goal would be to combine low cost onboard cameras with classification software on selected commercial flights to gather additional observations of early contrail development. The present study describes how the camera footage from the two campaigns was prepared for analysis and how the ambient conditions, such as time of day, haze, sun angle, limited the usable observations.

Deployment would require further assessment of camera installation, image quality and classifier performance across aircraft and operating conditions. Whether the resulting observations could contribute to data assimilation or improved meteorological forecasts or predictions on contrail persistence remains a question for future investigation.

---

2. The manuscript defines "short-lived contrail" and "persistent" categories not as observed in the imagery but via proxy of chase-plane measured ambient conditions of temperature, SAc and RHi > 100%. This definition glosses over the current open research question about agreement levels between SAc/measured RHi and _observed_ contrails (most recently embodied by Low et al 2025 that is cited elsewhere in this manuscript and https://egusphere.copernicus.org/preprints/2026/egusphere-2026-1171/). Authors do not attempt to visually verify in the imagery (even in a sample) the agreement level between the chase plane SAC/RHi measurements and the camera observations of contrail formation/persistence, and this represents a missed opportunity to improve scientific significance with this manuscript.

We have included a visual verification of the SAC measurements in Section 3 to investigate their agreement. We present a confusion matrix with the outcome and example images.

Section 3.

The agreement between SAc and manual visual observation was investigated and summarized in Table 9 for ND-MAX. Taking the camera-based manual labels as the observational reference, SAc

correctly indicates contrail formation in 97.9% of frames where a contrail is visible and no formation in 71.1% of frames where none is visible. Because contrails appear in about 73% of frames, κ is the more informative. Disagreement is asymmetric since about 84\% of the estimated 1,525 disagreeing frames are cases where SAc predicts formation but no contrail is visible (false-alarm ratio of 9.8%), compared with 248 frames (2.1% of observed contrails) where a contrail is seen but SAC predicts none.

SAC is a condition for formation, not for visibility, so frames with marginal ΔT_SAC may produce contrails too thin or short-lived to be seen by the camera. Other possible contributors are ambient conditions such as glare or haze, timing or position differences between the chase and test aircraft, and the fixed engine and fuel assumptions in SAC. Some examples are shown in Fig.7. These results put a ceiling on how well the image-based classifier trained on these SAC labels can score.

---

3. The definition of "persistent" being only RHi > 100% is misleadingly different from most other recent papers in the field which define persistent contrails as confirmed to be existing for at least several minutes (typically 10+ minutes). The vast majority of contrails formed are short-lived in the sense of having little to no climate impact because they exist only as linear contrails for seconds or only a few minutes, not persisting long enough to spread into contrail cirrus and have a radiative forcing effect on a large area. Even for a rear-pointing camera with good viewing angle and otherwise clear conditions, my napkin math indicates it's very unlikely the images could contain evidence of the aircraft's contrail persisting/not beyond 60 seconds from formation, when the aircraft is traveling at cruising speeds. Many recent papers have asserted the case for prioritizing forecasting the climate-relevant "big hit" contrails (very long-lived contrails, sometimes 12+ hours), and the burden of proof is on the authors to justify the claim that camera observations of SAc/RHi for at most 60 seconds would be useful in forecasting climate-relevant persistent contrail formation.

The field of view and the reliance on SAC/RHi for training are indeed limitations of this methodology. Those are made more clear in several places in the manuscript. As discussed in the first comment's response, the phrasing is being corrected to suggest the use of the proposed method as additional inexpensive real-time information on probable contrail formation that may be used with other methods of forecasting and validation.

Section III (Results)

While ice-supersaturation increases the likelihood of a persistent contrail, the level of correlation is still a question for discussion and the most common definition for persistence involves exhibiting a lifetime above a threshold value. This study sorely employs thermodynamic properties

and is therefore detecting contrail presence and presence of ice-supersaturated regions. However, for the classification aspect of this study, a contrail is labeled persistent by passing the SAC and the ice-supersaturation criterion and short-lived if it passes the SAC. This assumption is used for simplification and demonstration purposes and it merely indicates a higher likelihood for each contrail.

Section IV (Discussion)

As explained in earlier sections, for the purposes of this study the classification of contrails was simplified into three classes: absent (SAC not satisfied), short-lived (SAC satisfied but not the ice-supersaturation criterion), and persistent (both satisfied). In practice, these correspond to predictions of contrail absence, contrail presence, and contrail presence under ice-supersaturated conditions, respectively. Integration with satellite imagery to determine the duration of the contrail would help establish the actual likelihood of persistence.

---

4. Potentially important dataset details not reported or analyzed:

- At which times was the chase plane inside the contrail/jet-vortex of the lead plane (measuring RHi that is strongly influenced by the jet exhaust), and when was it outside the contrail/vortex (measuring purely ambient conditions)?
- Which specific humidity data streams from the campaigns were used by the authors, whether they suffer from thermal lag, and if so how that might affect the label noise in the Authors' dataset?

Thank you for the feedback. We know that the humidity data in both campaigns were calculated from other measured thermodynamic properties but we have been unable to receive further information to address the above questions. However, we have mentioned the possible impact of the lead's wake in Section 4.

---

5. Claiming causality when reporting correlations:

- The authors report the reason for why F1 scores vary across the categories and campaigns is due to camera placement/orientation issues, and it may be so; but no error bars are estimated, and no evidence or statistical/causal analysis is presented to justify this particular cause compared with any other potential cause (such as spurious correlations in the datasets, or propulsion efficiency of one campaign's aircraft being closer to the modeled value of 0.33 compared to the other campaign's aircraft, etc).

At this point in time, the limited test flights and camera placements steered us to a qualitative recommendation of camera placements, as opposed to a statistical/causal analysis. The phrasing was changed by highlighting other possible causes, including differences in measurements, propulsion, video durations and number of flight days, and image extraction intervals. We mention the measurement acquisition may have differed, but we do not have access to more specific details.

Section IV Discussion:

The results indicate the potential of airborne cameras for detecting early contrail formation and ice supersaturation. As explained in earlier sections, for the purposes of this study the classification of contrails was simplified into three classes: absent (SAC not satisfied), short-lived (SAC satisfied but not the ice-supersaturation criterion), and persistent (both satisfied). In practice, these correspond to predictions of contrail absence, contrail presence, and contrail presence under ice-supersaturated conditions, respectively. Integration with satellite imagery to determine the duration of the contrail would help establish the actual likelihood of persistence.

For the ND-MAX campaign, the absent class was predicted most accurately, followed by the short-lived class. The persistent class was the most difficult to separate, with most misclassifications assigned to the short-lived class. This is not surprising, since the determined ice-supersaturated regions can be small, as seen in Fig.~\ref{sac2} and therefore challenging to capture with indirect timestamps, the response rate of the instrumentation and the chase-aircraft sampling configuration.

The trend is inverted for the ecoDemonstrator campaign, where the persistent class performed best. One possible explanation for the inversion is the sensitivity of the observations to ice crystal number. The ecoDemonstrator's lean-burn combustor produces lower nvPM emissions, and therefore also fewer and larger ice crystals.

In turn, that makes the contrail brighter in ice supersaturated air conditions, while limiting crystal growth in subsaturated ones. At higher ice crystal numbers, the available water vapor is shared among more particles, potentially weakening this contrast. This is a hypothesis that would require testing with modeled ice crystal numbers and optical parameters extracted from the observations.

The contrasting results of the two campaigns may be attributed to several factors, none of which can be isolated with the present data. As discussed above, the differences in aircraft and engines are one significant possible factor. The amount of days and hours of data from each campaign varies, as well, with only 4 days from ecoDemonstrator footage compared to 7 days from ND-MAX. Interestingly despite the additional data, similar class counts for persistence and 2 versus 5 second image intervals, ecoDemonstrator performed better in persistence. Differences in measurement acquisition may also contribute.

Camera orientation is another possible factor. The ecoDemonstrator camera captured only a short segment of the exhaust plume close to the engine. Since the distance behind the engine at which contrails form varies, some contrails could have been outside the camera's FOV. Fig. 6 shows example images from the two campaigns. In the ND-MAX footage, the aircraft's tail obstructs part

of the contrail at high roll angles because of the camera placement. The ecoDemonstrator camera is mounted lower and angled outward, which avoids this issue. In addition, a significant part of the fuselage and wing occupies the ND-MAX frame, whereas the outward orientation of the ecoDemonstrator camera minimizes intrusion from the airframe. However, the ecoDemonstrator view shows a considerably shorter segment of the contrail, whereas the ND-MAX footage extends several kilometers behind the aircraft in favorable conditions, and the length of a contrail can make it easier to distinguish from natural clouds. A higher or lower camera placement may prevent the tail from hiding the contrail at high roll angles. The loss of images to sun glare might be reduced by placing the camera high and pointing it downward, or by tinting the upper half of the FOV. An additional camera on the opposite side may help capture cases in which glare is too intense in one side's footage. Because the campaigns differ in several respects at once, these are suggestions for future campaigns rather than findings of this study.

---

6. Machine learning best practices not applied:

- The ML model (autoencoder pretrained on the images, then separately PCA on latent vector followed by 2-layer fully-connected network) is clearly described, but the description is overly verbose for how routine this type of single-image classification task has become, while at the same time failing to justify its choices. Fine-tuning a deep CNN that was trained on ImageNet so it can be domain-adapted to a relatively small classification dataset like this one has been a baseline tutorial-level ML technique for over a decade now (https://research.google/blog/train-your-own-image-classifier-with-inception-in-tensorflow/) and more recently agent-driven optimization of model performance metrics has become possible with open libraries like https://github.com/karpathy/autoresearch and https://github.com/google-research/era ; if these were suitably used the reader would have some reasonable confidence that performance metrics were being maximized among at least a few experimentally attempted alternatives. And again, the more scientifically significant finding would be directly related to the end goal of contrail forecasting or weather re-analysis improvements, for which scores on this classification task are an (as-yet unvalidated) proxy metric.
- It's not fully specified in the manuscript how video frames are assigned to cross validation folds. If they were selected randomly for k-fold cross validation, it can potentially cause over-reporting of metrics. This is because adjacent frames can be extremely similar to each other but end up spanning the cross validation boundary (effectively training on the test set). Using frames restricted to one flight-leg in the campaign as one fold in the k-fold scheme would yield more representative metrics of how the model will perform when applied to new imagery from new flights in the future. See for example https://nsojournals.onlinelibrary.wiley.com/doi/10.1111/ecog.02881

Based on the feedback on ML practices, we reduced descriptions of the trivial techniques and steered towards sufficient justification of modeling choices.

Section II (Methodology)

Frames of size 1920 x 1080 RGB from the NDMAX and 1024 x 768 RGB from the EcoDemonstrator are encoded to 512-dimensional embeddings by a convolutional residual autoencoder. It is trained with a combined loss of MSE/L1 / SSIM and weights of 0.55/0.15/0.30, empirically determined with a preliminary sweep. Embeddings are reduced by PCA to 10 components and classified by a small fully connected network into absent, short-lived and persistent classes, whose labels are derived from SAC applied to the measured atmospheric data. Architecture and training hyperparameters are given in Table 2. The primary classifier inputs are the PCA components. The model uses supplementary data, altitude, true heading, and sun angle, to train it to remain invariant to changes in the background and lighting conditions across images.

A reconstruction-trained autoencoder was used rather than fine-tuning an ImageNet-pretrained network, initially motivated by two considerations. The embeddings are label-independent and reusable for other target definitions. Additionally, each campaign provides only 4 to 7 flight days, and the frames within a day share commonalities. The labels are also not from camera observation but thermodynamic criteria. Therefore, we hoped to test a frozen encoder with a small classifier expecting it to be less likely than full fine-tuning to memorize the appearance of individual days.

The 512-dimensional embedding size was chosen for reconstruction and is still a large size for the classifier. Since the classifier is trained on several correlated frames, PCA is used to reduce the size into a more compact, decorrelated input. For the PCA, 10 principal components (96\% of dataset variance) were chosen. Above 10, the accuracy did not improve significantly, while reducing below 10 degraded it.

The classification neural network is created to categorize contrail status into three classes. It uses the SAC results to train on the reduced images. The model uses random 5 k-fold cross-validation. A weighting system is used for class imbalance, since not all three classes are present in the same number of frames. The weights are normalized to sum to the number of classes. The results presented in Sect. 3 compare the actual versus predicted classes in every contrail category for each image in a confusion matrix. Additional training information is summarized in Table 2.

Section III (Results)

The image frames were resized to 320x320 pixels to reduce the input vector to a standardized size that can be processed by an onboard computer. The autoencoder's training metrics shown in Table 4 suggest stable convergence with training loss below 0.002, and high reconstruction quality with PSNR above 37~dB and SSIM close to 1. The resulting latent space was normalized before entering the PCA to encourage more pronounced clustering.

The classification was performed using an ANN evaluated with k-fold cross validation to improve the robustness of the evaluation and reduce the risk of overfitting. Out-of-fold predictions are pooled for each campaign. Table 5 summarizes the number of images per class and the weight assigned to them to avoid an imbalanced dataset. The final per-class metrics and confusion matrix were computed over the complete set of out-of-fold predictions across all samples.

The results from the submitted manuscript have been kept as is, however, we have noted the temporal blocking k-fold cross-validation recommendation. We applied group per day k-folds to validate the model against an unseen day, and observed an F1 score of 57.7% for ND-MAX. The range of ambient and lighting conditions across images is likely increasing the difficulty of training when the model hasn't seen as many conditions. Increasing the amount of days of footage might improve accuracy. Unfortunately, we are limited to the data we have used. Other approaches for future work, also mentioned in Section 4, could include more techniques and architectures to be explored, such as image editing/filtering, .

Section IV (Discussion)

The range of ambient and lighting conditions across images challenges training for both campaigns. Hazy conditions and passage through clouds can significantly reduce the visibility of contrails. In addition, the autoencoder encodes the entire image, including irrelevant characteristics, so that all color and texture differences between images are encoded and can distract from contrail identification. Differences in sky color and background clouds can have the same effect. When the aircraft rolls, the surroundings are tilted and the sun position may change the lighting. The time of day changes the color of the sky, and different levels of haziness are observed. Moreover, the wake of the lead aircraft may skew the chase aircraft's measurements of thermodynamic properties. These are currently limitations for the model's improvement.

Limited image editing was tested and could be explored further as a means of improving model performance. Other approaches, such as ImageNet pretrained networks and segmentation algorithms, merit testing and comparison against the current autoencoder, PCA, and neural network architecture.

7. Citation support for claims is frequently low quality, a few selected examples are below (but, I strongly recommend all citations need to be audited manually by the human authors of this manuscript):

- Line 56: citation needed
- Line 76: Akhtar Martinez et al is an inappropriate citation for this claim about open-source availability. The "open-source availability" that actually democratized CoCiP and led to its recent "broad adoption" is Shapiro et al 2023 ("pycontrails: Python library for modeling aviation climate impacts").

- Line 110: the paper describing "integrating satellite imagery with meteorological and flight trajectory" is actually https://www.nature.com/articles/s44172-024-00329-7#Sec2
- Table 2: The overall propulsion efficiency (eta) parameter is set to 0.33 for both campaigns' data, as shown in Table 2 without citation; the aircraft in these two campaigns may have a different propulsion efficiency.
- Line 238: Citing Chatterjee and Choudhury 2025 is a strange choice. It would be better to cite the foundational reference for PCA (first reported in the year 1901) or a reference that provides specific precedence/justification for the choice of pretraining an autoencoder, applying PCA, then training another separate classification network on top of the PCA components.
- Line 275: citation needed (CoCiP *does* use SAC…)
- Line 340: 50% accuracy claim is not supported by this citation. There is indeed an erroneous claim made in that IATA report that "Recent trials were successful in identifying the detected cloud as an aircraft-generated cloud 50% of the time when using an artificial intelligence (AI) algorithm, and 80% of the time when the matching was done or assisted by a person [17]" but citation 17 in that IATA report is Geraedts et al 2023, and in Geraedts et al 2023 the actually peer-reviewed scientific claim is "Only 50% of the flights labeled as matches by the automated matching were labeled as matches by the humans" which is a claim about *contrail-to-flight attribution* and NOT about contrail *detection*. Ng et al 2023 (that this work already cites, but elsewhere) is one option for a suitable reference for accuracy of contrail detection algorithms on satellite imagery.

We have additionally performed a full manual audit of every citation in the manuscript, as recommended. Regarding the citations that were flagged by the reviewer:

- Line 56: We added the citation by Teoh (2022)

Section Introduction

For example, Teoh et al. (2022) found that strongly warming and cooling contrails over the North Atlantic were associated with differences in season, time of day, background cloud conditions and aircraft emissions.

- Line 76: The open-source implementation of CoCiP that underlies its recent broad adoption is pycontrails (Shapiro et al., 2023, doi:10.5281/zenodo.7776686); We had included it in the very first draft but chose to remove it since it was the Github repo and not a paper. We now cite it there.

Due to the fact that mixing and bulk cloud processes are represented within CoCiP using semi-analytical formulations, it remains more computationally efficient, and its open-source availability

through the pycontrails library (Shapiro, 2023) has facilitated its broad adoption across the research community.

- Line 110: Included the Sonabend-W reference for the integration of satellite imagery with meteorological and flight trajectory.

Sonabend-W et al (2024) used a machine learning model that was trained on contrail detection and collocated numerical weather data to forecast zones where contrail formation was likely. Each flight segment in the training dataset was labeled as matching or not matching a contrail detected in GOES-16 imagery, using wind advection of the flight trajectories. These predictions were used with CoCiP forecasts to select flights for a randomized crossover trial of contrail avoidance in commercial aviation. To test the feasibility of contrail avoidance for a single airline, pilots from American Airlines participated and made altitude adjustments based on updated machine learning forecasts and blinded human evaluators used GOES-16 imagery to assess whether individual flights produced detectable contrails.

- Table 2: Bräuer et al. calculates η=0.31 for the ND-MAX campaign (https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020GL092166), while Dischl et all use 0.33 in their model for any aircraft that appear in the region https://www.mdpi.com/2226-4310/9/9/485 . For the ND-MAX, the value was therefore changed to 0.31 and for the ecoDemonstrator to 0.415, gathered from the ICAO Aircraft Engine Emissions Databank.

- Line 238: Cited Pearson (1901) for PCA. We will also improve our explanation on the revised paper.

The footage is then discretized into image frames to be conveyed through an autoencoder, followed by Principal Component Analysis (PCA) (Pearson, 1901) for further dimensionality reduction.

- Line 275: Our wording here was wrong and contradicting our earlier statement in line 216. We are citing the 2012 paper by Schumann.

To model the full lifecycle of contrails, from initial formation conditions till dispersion, physics-based simulation tools such as the Contrail Cirrus Prediction model (CoCiP; Schumann, 2012) and the Aircraft Plume Chemistry, Emissions, and Microphysics Model (APCEMM; Fritz et al., 2020) have been developed throughout the years.

- Line 340: We are removing this citation. When citing the IATA report, there was also confusion among the authors what the report meant.