

We thank the reviewer for their further efforts in ensuring the correctness of our work. We believe that we have addressed all of the issues raised.

Thanks to the authors for their thoughtful and productive response, and the changes made. The revised manuscript offers a significant improvement over the original, and the changes made have substantially improved the clarity, making for an enjoyable read.

Two previously raised issues persist, which can be rectified with language changes alone:

RC3: The manuscript still implies the independence of the ensemble control and operational ERA5 data. These datasets have the same underlying model, so will exhibit shared systematic biases. The appendix analysis only establishes the distribution of the errors is similar, but doesn't establish their independence.

Hence, the benchmark is not representative of how different attribution algorithms perform when faced with errors in advection-critical quantities in ERA5 data.

This only needs acknowledging in relevant portions of the text, specifically:

- Sect. 2.1.2
- Lines 118-119 (where that advection error characteristics are said to be realistic - flaws of contrail model and of the meteorological data upstream are also baked in)
- Line 238: "real weather" minimises these errors.

We have audited the manuscript to ensure that the claims about advection error of ERA5 vs. the synthetic contrails being similar to that of ERA5 vs. real contrail detections mention these potential biases and do not overstate their realism.

This feature of the benchmarking is important for this work, because the character of advection error is the basis for CoAtSaC, so a simpler advection error could feasibly (though not necessarily) suit it particularly well, and cause it to benchmark better. Therefore, this limitation should also be emphasised with the presentation of relative benchmarks in Sect. 4.1 (around line 614).

Done

Along similar lines, the limitation of attribution when using pixel values (line 286) ought to be emphasised in section 4.1 as it is relevant for comparing these algorithms.

Done

RC13: The revised version still doesn't accurately reflect the approach of the modelling studies (in line 23). This is representative of some flaws in the logical flow of the introduction.

The model studies do scale the humidity in order to correct the frequency of ISSR occurrence and the RHi distribution within them. Therefore, it is not true that the studies "do not account for inaccuracies therein" (line 28).

However, it is true that the model studies don't use a humidity field that is correct in terms of e.g. the spatial distribution of ISSRs. Therefore, they establish a representative population of contrails (which are a suitable basis for their conclusions as listed in the manuscript), but fail in predicting the occurrence of specific individual contrails (which would be required to compare observations and model output).

As a result, the statement "the benefits achievable by performing contrail avoidance with existing weather forecasts by differ from those quantified" doesn't follow - making these conclusions doesn't need individual matches. Further, line 32 seems to imply that these issues could be resolved with observations, but it is the deficiency in weather data that restricts model/observation comparisons. Perhaps observations would be better presented as an alternative impact assessment method, as long as attributions to flights can be made.

We have reworked this section to focus less on evaluation and improvement of prediction models, which is not the direct subject of this study, and instead focus on using observations to evaluate operational contrail avoidance. We believe that the reworked text should address the concerns raised here.

And some technical points/clarifications:

- Line 90: "for comparing a contrail forecast model to satellite observations". Perhaps change this to "for comparing properties of generating aircraft to satellite observations", to avoid the need to

relitigate whether a contrail forecast model produces a population that is comparable on an individual basis.

We reworded to “comparing the per-flight predictions of a contrail forecast model to satellite observations, for the purposes of assessing the model’s utility for operational avoidance.”

- Line 161: 2 waypoints per pixel doesn't follow immediately unless a value is stated for C_Tflight. It would be clearest to state it here as well as having it in the table.

Done

- Line 221: Clarify if the smoothing imposed is a spatial Gaussian blur (if so, standard deviation of 1 in which distance units), or a Gaussian noise in opacity (in which case, s.d. of 1 seems large, is this justified?)

Clarified that this is a spatial Gaussian blur with a standard deviation of 1 pixel.

- Sect. 4.2.2: This section, moved from an appendix, is a little dense (especially lines 660–694). Perhaps Fig. 13 could be removed. Furthermore, it is not clear from Fig. 12b that the altitude dependence is very strong, especially for the precision.

We have moved Fig 13 into an appendix along with all of the discussion about it, and replaced it with a short summary of the key findings. We also softened the language regarding the trends observed in Fig. 12b.