An algorithm to retrieve peroxyacetyl nitrate from AIRS

Response to reviewer 1

Joshua L. Laughner, Susan S. Kulawik, and Vivienne H. Payne

October 7, 2025

We thank the reviewer for their comments. In particular, their question about whether a stricter threshold for the PC-based filter would filter out the spurious signal seen in the original Australian Fires case led us to identify a way to correctly filter out those soundings. (It was not as simple as increasing the threshold, but was straightforward and appears to be effective.)

Below we respond to the individual comments. The reviewer's comments will be shown in red, our response in blue, and changes made to the paper are shown in black block quotes. Unless otherwise indicated, page and line numbers correspond to the original paper. Figures, tables, or equations referenced as "Rn" are numbered within this response; if these are used in the changes to the paper, they will be replaced with the proper number in the final paper. Figures, tables, and equations numbered normally in our responses refer to the numbers in the revised paper.

Major issues

However, while the technical implementation appears robust, I find the scientific contribution limited in its current form. The novelty primarily lies in extending the CrIS-based PAN retrieval strategy to AIRS, but the manuscript stops short of fully exploiting this opportunity. In particular, the analysis is largely confined to comparisons with CrIS over a few case studies. The broader potential of the AIRS PAN dataset to provide scientific insight remains underexplored. The scientific impact would be significantly enhanced if the authors included additional analyses, such a preliminary global climatology or seasonal cycle of AIRS PAN, assessments of long-term or interannual variability, or regional investigations beyond biomass burning plumes.

It is computationally expensive to generate a long timeseries of data with this retrieval. Therefore, it is only practical to carry out these sorts of analyses after the retrieval has been incorporated into an operational processing environment with sufficient computational resources to process larger batches of data. This analysis was restricted to that which was practical to process in our scientific computing facility. We have added a pair of sentences to the end of the introduction to explain this:

> "Due to the computational cost of this retrieval, our analysis focuses on a few days with significant variation in PAN from major fires in the US and Australia. This product will be incorporated in the operational TRopospheric Ozone and its Precurors from Earth System Sounding (TROPESS) data processing in the future (`https://disc.gsfc.nasa.gov/information/mission-project?keywords=tropess&title=TROPESS`, last accessed 11 Sep 2025), which will enable analysis on a longer timeseries of data."

The authors note that low, warm clouds over oceans can be misinterpreted as PAN. Yet, similar clouds exist over land (e.g., tropical forests). Could the authors clarify why this misinterpretation would be less problematic over land?

We have added a new paragraph with our hypotheses to the end of Sect. 3.2:

> "Our hypothesis is that the reason the AIRS retrieval is affected by the low, warm clouds and CrIS is not is due either the difference in spectral windows used between the retrievals (Fig. 1), the difference in radiance noise between the instruments, or a combination of the two. Further, our hypothesis for why land soundings are much less impacted than ocean soundings is that it is more difficult to distinguish a low, warm cloud from an underlying ocean surface than a land surface."

Why are these low, warm clouds an issue for AIRS but apparently not for CrIS? Are CrIS retrievals performed "above clouds"? If so, wouldn't that introduce a bias in retrieved PAN due to lack of surface contribution? More clarification on this aspect is needed.

CrIS retrievals are not performed specifically "above clouds," though, of course the effect of a cloud must be captured in the radiative transfer. The only difference between the CrIS and AIRS retrievals are the selection of spectral windows for the retrieval and minor details of the strategy table. Our hypothesis is included in the new paragraph quoted in response to the previous comment: this is likely due to some combination of the specific spectral windows used in the retrieval and the different noise characteristics of the instruments.

In Figure 3, some PAN features seen by AIRS (e.g., near 45°N, 145°W and 50°N, 130°W) are absent in CrIS. Are these retrievals cloud-contaminated? A short discussion of these discrepancies would improve interpretation.

We have inserted a few additional sentences discussing those discrepancies:

> "Similarly, in the plume around 50°N, AIRS sees enhanced X_{PAN} further west than CrIS (around 150° W) and more to the northwest of the state of Washington (near 50°N, 125°W). From the cloud properties shown in Fig. 4, these are also potential cases of erroneous impact from clouds."

The paper would benefit from a clearer and more detailed description of how the decision tree was designed, trained, and applied. Specifically, is the nearest CrIS PAN value used only during the training phase, or is it required systematically for each AIRS retrieval at the application stage?

We clarified that the CrIS X_{PAN} is only used in training in the discussion around Table 5:

> "Note that the CrIS X_{PAN} is not an input to the decision tree; it is used only in training. This permits the decision tree to be applied to AIRS soundings without a coincidence CrIS sounding."

If the quality filtering process requires CrIS data on a systematic basis (i.e., for each AIRS retrieval), then the utility of the AIRS PAN retrievals becomes restricted to the CrIS era (i.e., post-2012). This undermines one of the main potential advantages of using AIRS — the opportunity to generate a long-term PAN time series starting from 2002.

As described in the response to the previous comment, the filtering process does not rely on CrIS in this way and thus can be applied to AIRS data before CrIS was launched.

By tailoring the AIRS quality filtering strictly based on CrIS, there is a risk of overly aligning the two datasets. This may introduce biases or lead to the rejection of potentially valid AIRS PAN retrievals in cases where CrIS retrievals are biased, noisy, or simply absent. For instance, even over land, the quality filtering seems to restrict useful retrievals close to strong emission sources or fires.

In principle, we agree that it would be preferable to validate the AIRS product against an independent dataset, as Payne et al. (2022) did with the ATom flight campaign for the CrIS PAN product. However, the majority of the ATom profiles were flown over ocean. Given our need to filter out soundings impacted by the low clouds discussed above and the greater sounding-to-sounding noise seen in the AIRS product, this made both ATom and a similar campaign (HIPPO) difficult to use for our study. Further, other in situ profiles of PAN over land are sparse enough that, given the large sounding-to-sounding noise in our product, individual comparisons with the limited number of over-land profiles would not be meaningful. These issues are already described at the start of Sect. 3.3. However, since CrIS has been validated against in situ data, we believe that tying to CrIS is the best approach to mitigate biases.

The current implementation applies the AIRS AVKs to the CrIS retrievals to enable direct comparison. However, in my understanding of Rodgers, applying the AVKs from one instrument to retrievals from another is generally appropriate only when the second instrument has significantly higher vertical resolution and information content. In that case, it can reasonably serve as a "truth" profile.

While AKs are indeed commonly used to compare remotely sensed data against higher resolution in situ profiles, they can also be used to address different vertical sensitivities between two remote sensing instruments. This exact case is covered in Sect. 4.3 of (Rodgers and Connor, 2003). In our case, since both the AIRS and CrIS retrievals use the same a priori profile, using this profile as the "comparison ensemble," $\mathbf{x}_c$, from Rodgers and Connor (2003) is appropriate, as both retrievals will be optimal with respect to that profile.

However, both AIRS and CrIS PAN retrievals have limited vertical sensitivity, with DOFS that would typically be well below 1, indicating no vertical information.

The averaging kernels shown incorporate the pressure weighting function. Since each level is weighted by its contribution to the total column, that reduces the overall AKs. We have added additional panels to Fig. 15 that show the sum of the rows of the profile AK matrix, which are not weighted by the column operator, and so give a sense of the total sensitivity of each level of the retrieval to all levels of the true profile. We have also added a new figure (Fig. 16) that compares the overall DOFs between AIRS and CrIS. While the AIRS DOFs are lower than those of CrIS, they indicate there is sensitivity to column PAN, especially in scenes with large enhancements.

Fig. 12 shows that both AIRS and CrIS exhibit heterogeneous and situation-dependent vertical sensitivities (their AVKs diverge markedly when surface temperature decreases). Given this, the assumption that AIRS AVKs alone can transform CrIS data into something comparable is questionable. Ideally, a symmetric or "two-way" treatment accounting for both sets of AVKs would be required for this inter-comparison (yet this is practically challenging and still not guaranteed to yield equivalence in a formal sense).

Although the CrIS AKs do not show up explicitly in our Eq. (2), Eq. (26) in Sect. 4.3 of Rodgers and Connor (2003) shows that both instruments' AKs are present when we take the difference of the AIRS columns with respect to the CrIS columns adjusted with our Eq. (2), as the CrIS AKs are implicitly in $\hat{\mathbf{x}}_{\mathrm{CrIS}}$ (i.e., $\hat{\mathbf{x}}_2$ in Rodgers and Connor (2003)).

I find it unfortunate that the discussion and analysis of the AIRS PAN product is currently limited to land. Such limitation significantly reduces its utility in key applications, such as tracing fire plumes, where a large fraction of the signal occurs over oceans. In the case of the Australian bushfires, for example, nearly the entire plume over the ocean is lost.

Fortunately, when we were double checking on the effect of the filter for one of your other comments, we found a way to improve the PC-based filter and allow ocean data through, and then update the decision tree filter to also handle ocean soundings. We have updated our conclusions to reflect this, including that we now believe users can use the ocean data for large PAN plumes, with the potential to apply custom filtering when needed.

Lines 132-142: This section is difficult to follow without prior knowledge of the MUSES algorithm. I recommend expanding the explanation with more technical details to make it more self-contained and accessible to readers unfamiliar with previous TROPESS-related publications.

We have expanded the paragraph discussing the difference between the initial state and a priori constraint:

> "An important distinction within MUSES is the difference between the a priori (or constraint) state vector and the initial state vector. The former is $\mathbf{x}_a$ in Eq. (1) and is a mathematical constraint on the optimal state vector, the latter is the starting point of $\mathbf{x}$ before the first iteration of the Levenberg-Marquardt solver. **This distinction is important within MUSES because it is a multi-step retrieval. The strategy table, mentioned above, defines which elements of the state vector will be retrieved in each step and whether or not the retrieved state for step i becomes the initial state for step $i+1$. For example, the retrieval may begin with an H_2O profile taken from**

a meteorological reanalysis as both the initial guess and the a priori constraint. An early step in the retrieval can then retrieve a new H_2O profile which is more consistent with the observed radiances. This new H_2O profile can then be used as an initial state for later steps (whether or not those steps retrieve H_2O). This can be important for weak absorbers, such as PAN, which need the profiles of strong thermal IR absorbers to be accurate for the scene in question so that the relatively small absorption feature of the weak absorber can be identified. We note that, for a given step, the initial state and a priori constraint can be the same but do not need to be. For later steps of the retrieval, the initial state will have been set by earlier retrieval steps **(as in the example given with H_2O) but the a priori constraint will remain the same for all steps. Or, the a priori constraint may be chosen to be a relatively simple profile to avoid imposing undue assumptions, while the initial state may be chosen to reflect a better estimate of the atmospheric state in that location to attempt to minimize the number of steps needed by the solver.**"

For OSS, we have added a sentence summarizing the benefits and drawbacks of OSS:

> "This allows OSS to efficiently simulate the radiances a specific instrument would observe by reducing the number of monochromatic wavelengths that must be modeled for a given instrument channel, but means that OSS must be trained for each instrument used in a retrieval separately."

However, we do recommend that readers interested in the details of OSS read the Moncet papers rather than relying on our description.

Section 3.4: Although I understand that deriving uncertainty estimates for retrieved quantities from satellite measurements is challenging, I remain unconvinced by the authors' approach. The reported uncertainty value (0.5 ppb) is based solely on the difference in NESR between AIRS and CrIS. However, the uncertainty should realistically vary significantly with factors such as thermal contrast, cloud coverage, PAN abundance, and others.

Our conclusion that 0.5 ppb was a reasonable estimate of the uncertainty was not solely due to the NESR analysis. As we said in the first paragraph of Sect. 3.4, it is consistent with the analysis in the previous section that showed the spread in AIRS-CrIS correlation was approximately 0.5 ppb. To reinforce this point, we have moved the original Fig. 15 forward to be Fig. 14, showing that a similar 0.5 ppb spread is seen when individual AIRS and CrIS soundings are correlated. We have also added an acknowledgement that the individual sounding errors will vary, but that we believe the 0.5 ppb estimate to be a reasonable average for the AIRS PAN product:

> "While we expect the error of individual soundings to vary depending on the specific atmospheric and surface conditions for each sounding, we believe 0.5 ppb to be a reasonable estimate of the typical uncertainty in the AIRS X_{PAN} data."

Section 3.4: I find the discussion on vertical sensitivity rather brief. For example, what is the typical DOFS of the AIRS PAN retrievals in fire plume regions versus remote areas? How do these values compare to those from CrIS?

We have added a new figure (Fig. 16) comparing the DOFS of the two retrievals and a new paragraph comparing the DOFS values:

> "Figure 16 shows the overall degrees of freedom (DOF) of signal for both the AIRS and CrIS products in the 2020-09-11 US West Coast Fire scene. From Fig. 16 panels a and b, we can see that the DOFs for the CrIS PAN product are grouped around 1, indicating that there is essentially always enough information to retrieval a single piece of vertical information in the form of a column average. In contrast, Fig. 16 panels c and d show that the AIRS DOFs are lower (centered around ∼0.5) with a wider distribution. Greater AIRS X_{PAN} values do tend to be associated with greater DOFs. This implies that the AIRS product will retain influence from the prior, particularly in background conditions, but can detect sufficiently large PAN enhancements."

Minor comments

Lines 27-31: Do the authors have an estimate of what fraction of the total APNs signal in the retrievals corresponds specifically to PAN? Given its longer lifetime relative to other APNs, could one expect its share to increase in aged plumes or background air.

We have added a number of references pointing to PAN as typically comprising 75% to 90% of APNs, with a caveat about wildfire plumes:

> "However, PAN typically comprises the majority (75% to 90%) of APNs in both remote areas (Roberts et al., 1998, 2002; Wolfe et al., 2007; Fischer et al., 2014) and urban plumes (LaFranchi et al., 2009). The fraction may be lower in wildfire plumes; Peng et al. (2021) hypothesize that an unknown APN could explain discrepancies in NO_x/CO ratios between their observations and model."

Section 2.1 would benefit from more technical information about the AIRS instrument, especially in relation to its suitability for PAN retrieval (spectral resolution, radiometric noise characteristics (especially compared to CrIS), spatial sampling and footprint size).

We have added this information as Table 1.

Lines 150-153: The manuscript mentions a "global survey" sampling approach with TROPESS products. It would be important to clarify what proportion of soundings are included in the final products. For instance, is it 1 out of 2 soundings, 1 out of 3, etc.? This has implications for data representativity.

We added a sentence describing the thinning strategy:

> "The default survey strategy processes one sounding in each $x^{\circ} \times x^{\circ}$ box over land and one out of every four such boxes over ocean. For the current products, x is either 0.7° or 0.8°."

CCl4 is not mentioned in the strategy table (Table 2), yet it has notable absorption features in the thermal infrared that could affect PAN retrievals, especially the spectral baseline. Is CCl4 explicitly fitted in the retrieval process? If not, how is its temporal variability accounted for?

We have added an explaination of how CCl_4 is accounted for:

> "CCl_4 is not retrieved (Table 3) but is simulated in the radiative transfer as an interferent, using climatological profiles scaled by yearly scale factors derived from ground based observations. The base climatological profiles vary with latitude and longitude in 30° and 60° bins, respectively, and were developed from MOZART model output (Brasseur et al., 1998)."

Lines 183-186: Could the use of different a priori profiles across regions introduce discontinuities in the retrieved PAN abundances at regional boundaries?

In principle, yes. But this is not unique to this product; many remote sensing products use priors that vary in space and can introduce such discontinuities. Correct application of the averaging kernels and a priori constraints when comparing to other datasets will account for the influence of the prior.

Lines 211-213: Are the threshold criteria used for AIRS the same as for CrIS? If so, is this appropriate given the different instrument characteristics (spectral resolution, sensitivity, etc.)?

The criteria are not the same as for CrIS. In any case, these criteria are not what determine the filtering of the actual product; that is the machine learning model. We have clarified this:

> "The AIRS data shown in Fig. 3 are those soundings which pass **prototype** quality flags **chosen based on quality flags for other thermal retrievals**, including sufficiently small radiance residual, surface temperature > 265 K, cloud top pressure (as retrieved in our algorithm) below the tropopause, and the quality of the H_2O retrieval in step 4 of Table 3. **(Note that these quality flags were for prototyping purposes only, and are not those used in the final product.)**"

Line 232 ("the filtering approach failed..."): Could stricter filtering criteria resolve this issue?

Yes. In fact, when we checked the PC values to respond to this question, we realized that using PCs derived from a ~ 100 cm^{-1} window around the PAN feature were more reliable than using just radiances in the PAN retrieval windows, and that increasing the threshold from -10 to 0 for PC2 would address the clouds in the Australian case. We have adjusted our conclusions to reflect this, though we do still advise caution as the variation in the effective threshold between these two test cases suggests there may be more variation in the appropriate threshold across time and space.

Lines 302-304: These statements could benefit from clarification in the case of the Australian Bush Fires. Much of the plume appears to be missing over the ocean, and the soundings over land seem relatively noisy. The observation that AIRS shows no PAN enhancement, similarly to CrIS, should be interpreted with caution. The agreement between the two instruments in this case does not necessarily validate the accuracy of the AIRS retrievals, especially in light of the limited data coverage.

With the improved PC-based filter, we are now able to show in the new Fig. 10 that AIRS does see a PAN plume heading towards New Zealand.

In Fig. 8 (and similar figures), it is difficult to assess the differences in spatial sampling and resolution between the CrIS and AIRS soundings. Including a zoomed-in view might help better illustrate these differences.

The spatial sampling of the two instruments is not the point of these figures. We have added Table 1 that shows how the two instruments compare in that regard.

Lines 321–326: Could you clarify whether the intention is to recommend that the AIRS PAN product be used primarily as $10^\circ \times 10^\circ$ spatial averages? If so, this seems quite restrictive.

We have added a cross reference to Sect. 4 where we discuss this in more detail. In brief, for background concentrations of PAN, we recommend averaging approximately 140 soundings to reduce the random uncertainty sufficiently, but for strong plumes, individual soundings can be used.

Lines 327–330: Would it be feasible to implement a similar H_2O bias correction for the AIRS PAN retrievals as is done for CrIS?

The significant random uncertainty in the individual AIRS soundings makes doing so difficult. We have added text acknowledging this:

> "Payne et al. (2022) were able to derive the CrIS bias correction through comparison between CrIS and in situ background X_{PAN} values. The need to average a significant number of AIRS soundings to reduce the random sounding-to-sounding noise makes it difficult to identify any relationship between AIRS X_{PAN} values and H_2O column amounts."

Technical comments

All technical comments have been addressed, thank you for catching the typos.

References

Brasseur, G. P., Hauglustaine, D. A., Walters, S., Rasch, P. J., Müller, J., Granier, C., and Tie, X. X.: MOZART, a global chemical transport model for ozone and related chemical tracers: 1. Model description, Journal of Geophysical Research: Atmospheres, 103, 28 265–28 289, https://doi.org/10.1029/98jd02397, 1998.

Fischer, E. V., Jacob, D. J., Yantosca, R. M., Sulprizio, M. P., Millet, D. B., Mao, J., Paulot, F., Singh, H. B., Roiger, A., Ries, L., Talbot, R., Dzepina, K., and Pandey Deolal, S.: Atmospheric peroxyacetyl nitrate (PAN): a global budget and source attribution, Atmospheric Chemistry and Physics, 14, 2679–2698, https://doi.org/10.5194/acp-14-2679-2014, 2014.

LaFranchi, B. W., Wolfe, G. M., Thornton, J. A., Harrold, S. A., Browne, E. C., Min, K. E., Wooldridge, P. J., Gilman, J. B., Kuster, W. C., Goldan, P. D., de Gouw, J. A., McKay, M., Goldstein, A. H., Ren, X., Mao, J., and Cohen, R. C.: Closing the peroxy acetyl nitrate budget: observations of acyl peroxy nitrates (PAN, PPN, and MPAN) during BEARPEX 2007, Atmospheric Chemistry and Physics, 9, 7623–7641, https://doi.org/10.5194/acp-9-7623-2009, 2009.

Payne, V. H., Kulawik, S. S., Fischer, E. V., Brewer, J. F., Huey, L. G., Miyazaki, K., Worden, J. R., Bowman, K. W., Hintsa, E. J., Moore, F., Elkins, J. W., and Juncosa Calahorrano, J.: Satellite measurements of peroxyacetyl nitrate from the Cross-Track Infrared Sounder: comparison with ATom aircraft measurements, Atmos. Meas. Tech., 15, 3497–3511, https://doi.org/10.5194/amt-15-3497-2022, 2022.

Peng, Q., Palm, B. B., Fredrickson, C. D., Lee, B. H., Hall, S. R., Ullmann, K., Campos, T., Weinheimer, A. J., Apel, E. C., Flocke, F., Permar, W., Hu, L., Garofalo, L. A., Pothier, M. A., Farmer, D. K., Ku, I.-T., Sullivan, A. P., Collett, J. L., Fischer, E., and Thornton, J. A.: Observations and Modeling of NOx Photochemistry and Fate in Fresh Wildfire Plumes, ACS Earth and Space Chemistry, 5, 2652–2667, https://doi.org/10.1021/acsearthspacechem.1c00086, 2021.

Roberts, J. M., Williams, J., Baumann, K., Buhr, M. P., Goldan, P. D., Holloway, J., Hübler, G., Kuster, W. C., McKeen, S. A., Ryerson, T. B., Trainer, M., Williams, E. J., Fehsenfeld, F. C., Bertman, S. B., Nouaime, G., Seaver, C., Grodzinsky, G., Rodgers, M., and Young, V. L.: Measurements of PAN, PPN, and MPAN made during the 1994 and 1995 Nashville Intensives of the Southern Oxidant Study: Implications for regional ozone production from biogenic hydrocarbons, Journal of Geophysical Research: Atmospheres, 103, 22 473–22 490, https://doi.org/10.1029/98jd01637, 1998.

Roberts, J. M., Flocke, F., Stroud, C. A., Hereid, D., Williams, E., Fehsenfeld, F., Brune, W., Martinez, M., and Harder, H.: Ground-based measurements of peroxycarboxylic nitric anhydrides (PANs) during the 1999 Southern Oxidants Study Nashville Intensive, Journal of Geophysical Research: Atmospheres, 107, https://doi.org/10.1029/2001jd000947, 2002.

Rodgers, C. D. and Connor, B. J.: Intercomparison of remote sounding instruments, J. Geophys. Res. Atmos., 108, https://doi.org/10.1029/2002jd002299, 2003.

Wolfe, G. M., Thornton, J. A., McNeill, V. F., Jaffe, D. A., Reidmiller, D., Chand, D., Smith, J., Swartzendruber, P., Flocke, F., and Zheng, W.: Influence of trans-Pacific pollution transport on acyl peroxy nitrate abundances and speciation at Mount Bachelor Observatory during INTEX-B, Atmospheric Chemistry and Physics, 7, 5309–5325, https://doi.org/10.5194/acp-7-5309-2007, 2007.