Response to reviewers

We would like to thank the reviewers for their thoughtful and constructive review of this work. We feel that addressing the reviewers' questions and incorporating their useful suggestions has led to significant improvements in the manuscript.

Reviewer comments are in plain text below, while our responses are in bold/italic.

Anonymous Referee #1

Received and published: 11 July 2014

General Comments

The manuscript by Payne et al. titled "Satellite observations of peroxyacetyl nitrate from the Aura Tropospheric Emission Spectrometer" describes the development of a new satellite retrieval product based on TES observations. Routine satellite observations of PAN would enable better understanding of the nitrogen cycle; current satellite products for PAN are limited to limb-sounding instruments which can not see into the troposphere. The manuscript establishes the framework for a PAN product based on TES observations and analyzes a limited number of actual TES PAN retrievals.

Overall, the paper does a thorough job of describing the PAN optimal estimation-based retrieval algorithm and its various components. Results are presented both for simulated retrievals, where PAN concentrations are specified, and for one month of actual TES observations, where the atmospheric PAN concentrations are unknown. The task of developing a PAN product is challenging for several reasons, but mainly because of the weak radiative sensitivity. Rigorously validating the retrieval product is apparently not yet feasible because of the lack of in-situ data.

Since the paper does not report actual validation results, the validity of the algorithm must be judged from a rigorous analysis of possible retrieval errors and from qualitative observations, such as the 'reasonableness' of actual retrieval results in particular contexts. In these two areas, the manuscript should be improved. Specific suggestions are included below.

Specific Comments

p. 5351, l. 14. The meaning of '... close to the true state ...' is unclear. How close is close enough? Can this statement be made more quantitative?

The retrieved state needs to be close enough to the true state such that the Jacobian, dL/dx_hat is approximately equal to dL/dx. Further discussion of "close enough" can be found in Section 5.1 of Rodgers (2000). We did not attempt to put quantitative bounds on that statement for this work. This would be a non-trivial undertaking and we are not convinced that doing so would add anything substantive to the analysis.

p. 5352, l. 5. Does Sn only represent instrument noise, or does it also represent systematic radiance errors relative to the forward model (e.g., spectroscopic errors)? Are forward model errors represented somewhere else?

Yes, here Sn only represents instrument noise. The reviewer is right to point out that we did not include forward model errors in Equation 1. Strictly speaking, there should be an additional term, G*delta_f, where delta_f represents the error in the forward model relative to the real physics. Delta_f would include spectroscopic errors, amongst other things.

p. 5352, l. 8. What is meant by 'relatively linear'?

'Relatively linear" means that although the retrieval problem itself is non-linear (and requires iteration to reach a solution), a linearization about some prior state is adequate to find a solution. An in-depth discussion of retrieval linearity can be found in Rodgers (2000).

p. 5352, l. 24. Some justification should be given for assuming linearly varying surface emissivity. Are there at least some materials where surface emissivity in this spectral region are documented?

We thank the reviewer for raising this important point that we had neglected to discuss fully. We have now added a new figure (nominally Figure 8 - see below) and the discussion paragraph that follows.

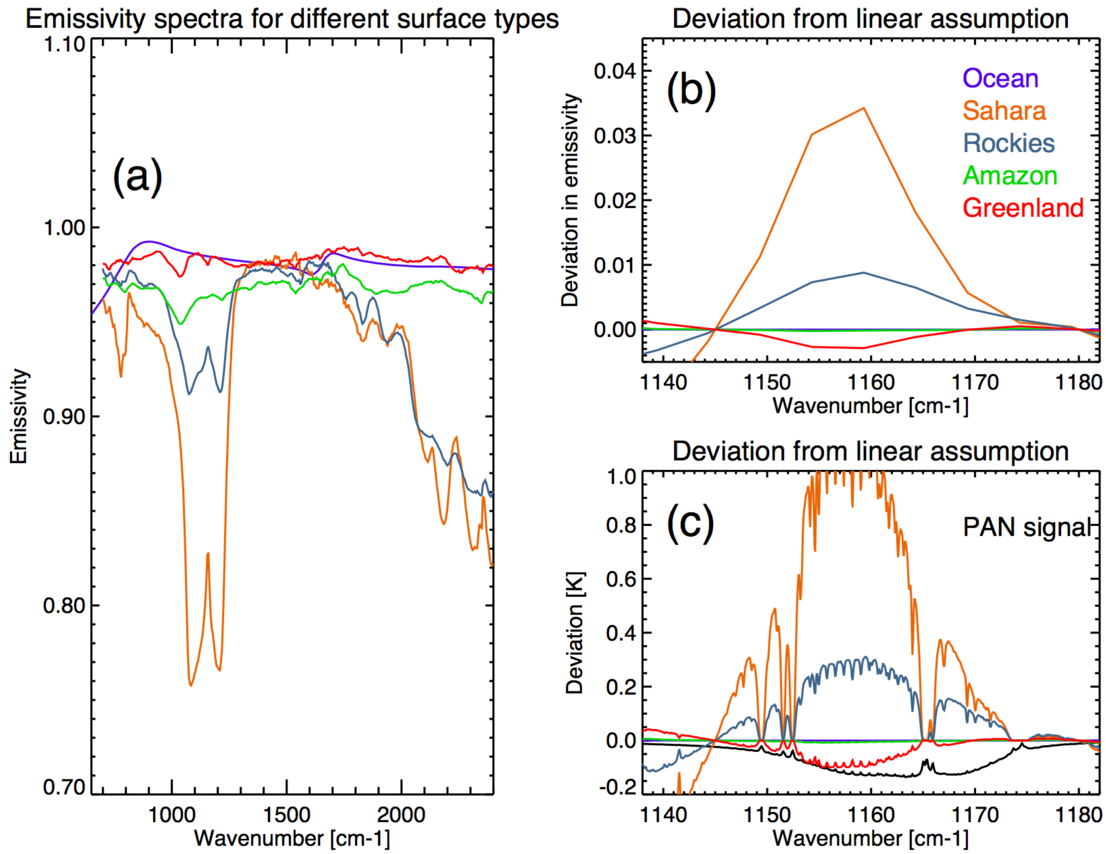

Figure 8: **(a) Emissivity spectra for different surface types, taken from the UW/CIMSS HSR emissivity database. (b) Deviation of the emissivity spectra from a linear extrapolation between 1145 and 1180 cm^{-1}. (c) Error in clear-sky brightness temperature resulting from the deviation of emissivity from the linear assumption. The black line in (c) represents the PAN signal associated with the July extratropical „maximum aloft“ profile shown in Figure 4(e) and is shown for the purpose of illustrating the spectral shape of the PAN feature relative to the spectral shape of the signal associated with deviation from linear surface emissivity.**

"If the surface emissivity and the cloud optical depth can be assumed to vary linearly with wavenumber across the ~40 cm^{-1} wide PAN spectral region, we can assume that the impact of uncertainties in these quantities are also made small by our choice of retrieval strategy. Over much of the Earth's surface, the assumption of linear variation in emissivity between 1140 and 1180 cm^{-1} is reasonable. However, this assumption is problematic over bare rocky or sandy surfaces, due to the silicate feature centered around 1160 cm^{-1}. Figure 8(a) shows emissivity spectra from the University of Wisconsin Co-operative Institute for Meteorological and Satellite Studies (UW/CIMSS) High Spectral Resolution (HSR) emissivity database (Borbas et al, 2007). A strong surface silicate feature will generally lead to a high initial χ^2 value (Section 3.2), and so retrievals are generally not attempted for such cases. If the retrieval were to be

attempted, the silicate feature would likely result in an underestimate of PAN (Figure 8 (b)). Spectral variation in the emissivity over snow or icy surfaces (see, for example, the "Greenland" plots in Figure 8) leads to an error in the opposite sense from the silicate feature. However, since the shape of the deviation from linear emissivity does not match the shape of the PAN feature, any significant deviation from the linear emissivity will tend to result in high χ^2 values, which either mean that the retrieval is not attempted in the first place, or that it will be rejected according to a final quality flag. (Retrievals where the final χ^2 is greater than 1.5 are flagged as "bad"). Evidence suggests that cases where a non-linear emissivity component produces a modeled radiance error that is equivalent in magnitude to a ~0.1 ppbv change in PAN would not pass quality control. We acknowledge that the variation in snow and sea-ice emissivity spectra is relatively large (Borbas and Ruston, 2010) and that it not completely outside the realm of possibility that there could exist surfaces for which the shape of the deviation from linear emissivity is a good match for the shape of the PAN signal. We will provide an upper bound of +0.1 ppbv "emissivity bias" for snow/ice surfaces and an equivalent bound of -0.1 ppbv emissivity bias for rocky surfaces.

In addition, we note that we did not find any sign of an increase in the final χ^2 values for the high latitude cases (or for the extreme high PAN cases) shown in the "Pacific Transport" map.

p. 5354, l. 7. For a given observation location, what are the actual criteria for selecting the a priori category?

For each model grid box, we classify the model profile in that grid box according to one of the six categories described above and use the relevant average profile for that category as the retrieval a priori. The text has been updated with this re-wording of the explanation.

p. 5354, l. 12. What is meant by 'entering null space'?

What was meant here was that if the initial guess is set too small, then we can enter a space where the Jacobians are effectively zero. (A vanishingly small perturbation of a vanishingly small number approaches zero.) If the Jacobians are effectively zero, then the retrieval can never move away from the initial guess. The text has been updated to include this information.

p. 5354, l. 18. Consider deleting 'relatively'

Done.

p. 5354, l. 19. Can this a priori variance value be interpreted as a percentage or fractional variability? How does this variance value compare to variances for other trace gases retrieved by TES?

We thank the reviewers for catching this. We should not have labeled this as S_a^{-1}, when

we ought to have labeled it as a constraint matrix, denoted by, say, R. The constraint matrix used here cannot be interpreted as an a priori covariance.

The main point here is that the retrieval is free to move. We wished to allow the retrieval freedom in both the shape and the magnitude of the PAN profile values. The constraint is very loose compared to the constraints for other TES trace gas retrievals, for which we tend to have better a priori knowledge of the distribution. (One exception to this is ammonia, where we have relatively poor prior knowledge of the global distribution and where the values can vary widely.)

We do anticipate re-examining both constraints and prior covariances pending further work in assessing the variability of the TES PAN product on global spatial scales and long-term time scales.

p. 5354, l. 21. I recommend revising this paragraph. In optimal estimation theory (i.e., Rodgers' book), the off-diagonal elements of the a priori covariance matrix can not be 'tuned', but really describe the expected or observed a priori correlations in trace gas concentrations at different levels. Setting the off-diagonal elements to 0 is equivalent to assuming a vanishingly small vertical correlation length.

As detailed above, we acknowledge that we should not have labeled the constraint matrix as S_a^{-1}***.***

Ideally, yes, we would have a good estimate of the expected or observed a priori correlations in trace gas concentrations at different levels. For PAN, we really don't have reliable prior information, certainly not in a global sense.

In this case, the correlations among the PAN jacobians throughout the troposphere effectively correlates the results at the different levels, even without off-diagonal constraints.

p. 5355, l. 13. Could some vertical information be theoretically possible using different microwindows, or is there a fundamental reason why this is unrealistic?

For the band that we are using, the shape of the PAN signal is not sensitive to the details of the vertical structure of the PAN profile. We had previously attempted to make this point using Figure 2. We have updated Figure 2 in an attempt to clarify this point. Figure 2 now shows three highly elevated PAN profiles, each with different vertical structure, alongside their associated brightness temperature signal. Also shown are the signals normalized to the maximum strongest signal of the three cases, in order to demonstrate that although the shapes of the profiles are different, the spectral shape of the resulting signal as seen by TES is the same. The updated figure is reproduced below:

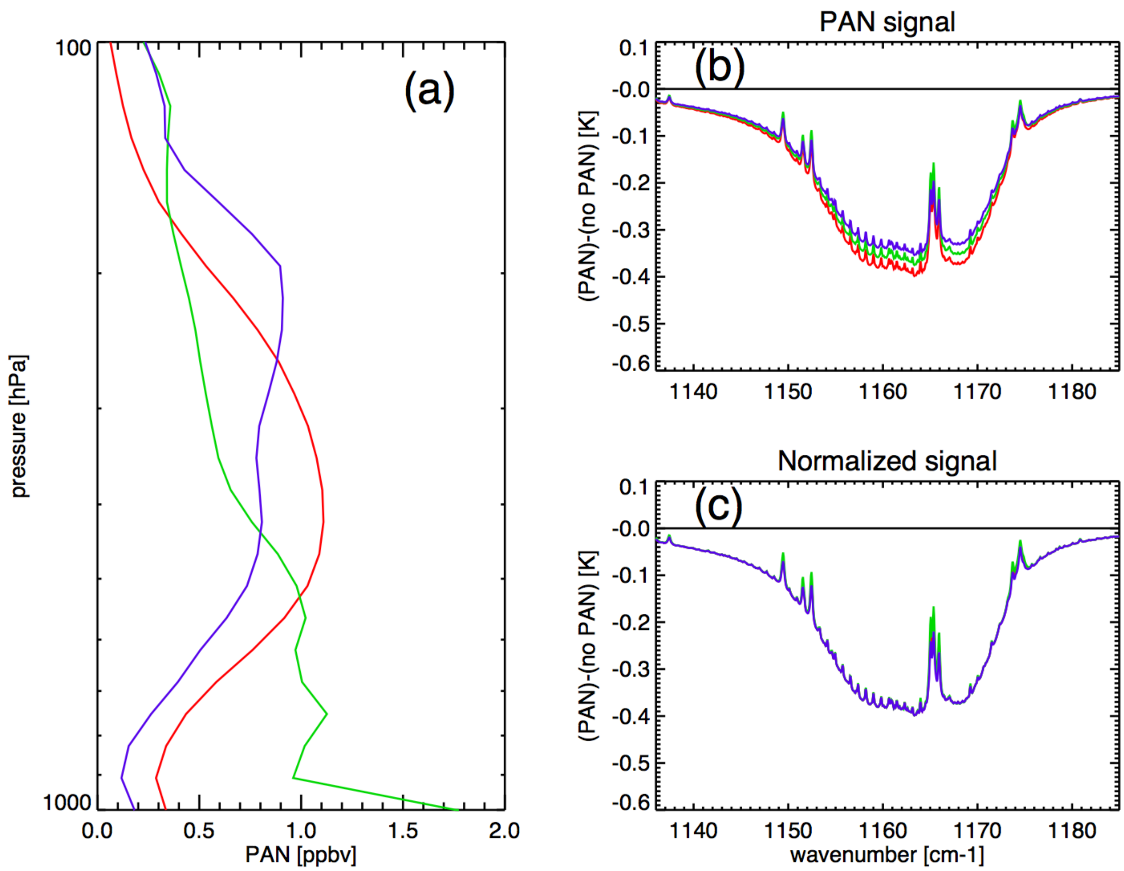

Figure 2: **(a) Three highly elevated PAN profiles, each with different vertical structure. (b) Brightness temperature signals resulting from the three elevated profiles. (c) As (b), but with each the signals normalized to the greatest brightness temperature difference value shown in (b)**

(For cases where the true PAN profile exhibits a maximum in the boundary layer, the extent to which we are sensitive to near-surface PAN is sensitive to the surface temperature and the thermal contrast between the surface and the lowermost atmosphere.)

The reviewer may also be asking about the possibility of using other PAN absorption bands. We have attempted to address this issue in our response to Reviewer 2's question about other PAN absorption bands in the thermal IR. In short, none of the other thermal IR PAN bands are viable possibilities for use in TES retrievals.

p. 5356, l. 12. Here it would be helpful to briefly review the physical basis of the effects of clouds on TES retrievals.

We have added the following text: "In the TES operational algorithm, trace gas retrievals are routinely performed in the presence of clouds. The clouds are implemented in the forward model as a single layer Gaussian vertical profile parameterized by a cloud height and a set of frequency-dependent effective (non-

scattering) optical depths (Kulawik et al., 2006). Cloud optical depths are not expected to show spectral structure within the ~40 cm^{-1} range of the PAN band used in this work."

p. 5357, Section 3.6. For the retrieval simulations, instead of comparing x_rtv with x_true, it would be more informative to compare x_rtv with the quantity x_a + A(x_true - x_a), as described in Rodgers' book, since this would account for the influence of the a priori.

We agree. We have updated Figure 7 (updated figure shown below in response to the next comment) to show this quantity and updated the text where Figure 7 is discussed.

p. 5357, l. 25. When evaluating the RMS differences between the true and retrieved values for the simulations, these differences should be compared to the RMS differences between the true and a priori values. If the retrieval algorithm has skill, the true/retrieved RMS differences should be substantially smaller than the true/a priori RMS differences.

Figure 7 has been updated to show the RMS differences between the true and retrieved values for the simulations alongside the RMS differences between the true and a priori values. The updated Figure is shown here below:

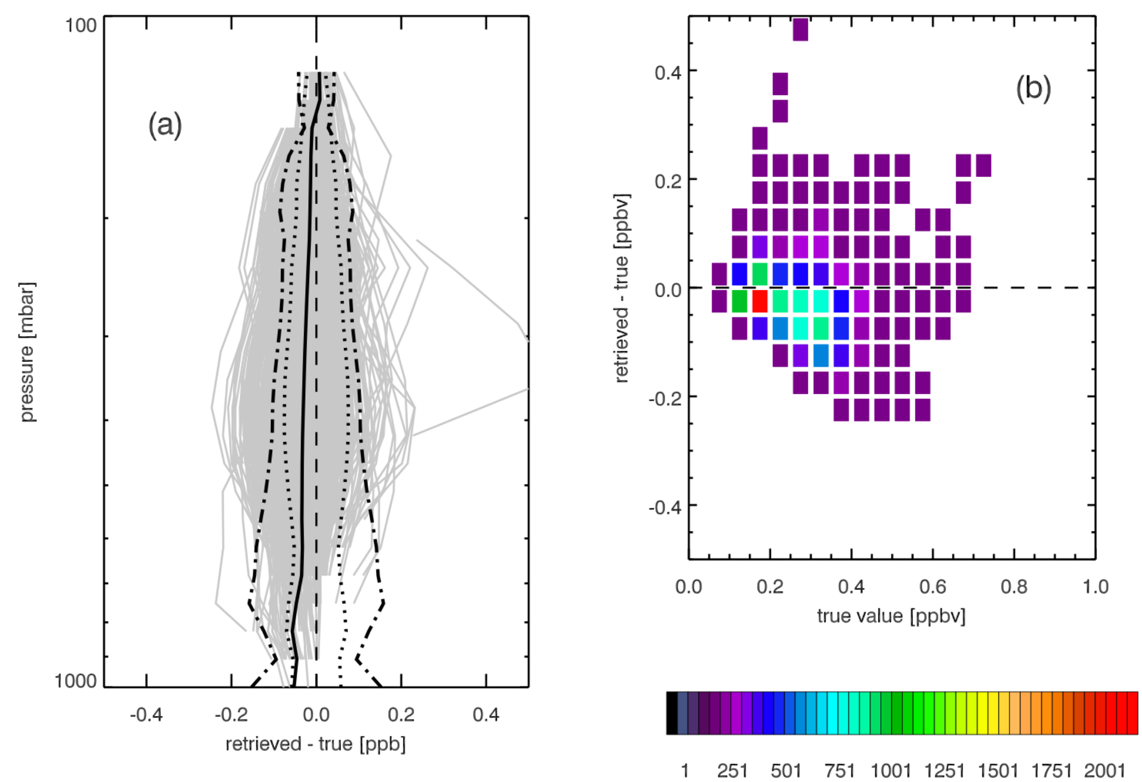

We have added the following text to the discussion of this Figure:

" We see that the RMS difference between the true and retrieved states is less than the RMS difference between the true and prior states for all altitudes,

showing that the retrievals have skill. The retrievals appear to show greatest skill between ~400 and 800 hPa."

p. 5358, l. 11. Is it known what surface types (e.g., water, vegetation, etc.) exhibit surface emissivity which varies linearly over the PAN microwindows? Ideally, simulations should be performed where modeled radiances are based on realistic surface emissivity data. Why is this source of retrieval error not included in Table 2?

This is an important point and we agree that it merits further discussion in the paper. Yes, it is known what surface types exhibit surface emissivity that varies linearly over the PAN microwindows. As the reviewer suggests, water and vegetated surfaces exhibit surface emissivity that varies linearly over the PAN spectral range used here. Sandy/desert surfaces, and to a lesser extent bare rocky surfaces, are known to be problematic due to the presence of the silicate feature centered around 1160 cm^{-1}. The shape of the silicate feature acts in the opposite direction to the signal of an increase in PAN. Retrievals over desert regions are generally never even attempted. We reject these cases due to high initial chi-squared values. Retrievals over rocky surfaces are generally not attempted due to initial chi-squared values. Snow covered surfaces can also exhibit non-linear spectral variation in the 1140-1180 cm^{-1} range. However, in general we find that significant deviation from the linear emissivity assumption (where "significant" is a deviation that would lead to a radiance signature comparable to the PAN signal) leads to poor spectral fits. Under these circumstances, the retrieval is either never attempted or will be flagged as "bad", based on the final chi-squared value.

In response to a number of reviewer comments, we have now added a number of error terms, including emissivity, O3, N2O, CFC-12 and surface temperature error estimates, to Table 2 (see below).

p. 5359, l. 16. The potential effects of O3 and N2O on the PAN retrievals are not quantitatively investigated. For both gases, it is not clear that the variability of the absorption features within the PAN microwindows would somehow prevent retrieval bias. Ideally, this section should include simulations where the O3 and N2O profiles are varied (relative to the assumed profiles in the retrieval), in roughly the same manner that H2O biases were investigated.

We have now quantitatively investigated these errors.

For O3, we assume a 40 % RMS uncertainty (based on TES O3 validation against ozonesondes at mid-latitudes – see Nassar et al. (2008)). We find that this RMS ozone uncertainty results in a ~0.1 ppbv RMS uncertainty in the PAN retrievals.

For N2O, we assume a 2% RMS uncertainty (based on global variability in tropospheric N2O VMR, as observed from HIPPO aircraft measurements). We find that this RMS N2O uncertainty results in a ~0.05 ppbv RM uncertainty in the PAN retrievals.

p. 5360, l. 2. What is the aggregate uncertainty from all of the sources of retrieval uncertainty?

Please see below for a revised version of Table 2, with an expanded list of uncertainty terms for a profile with 0.5 ppbv of PAN in the mid-troposphere. In the calculation of the aggregate uncertainty, we choose not to include the terms associated with CFC-12 (since this is largely correctable in future algorithm versions) or with the problematic surface emissivity types. Summing the squares of independent error terms 1-8 in the Table gives an aggregate uncertainty of 43 %.

With the exception of the spectroscopic error, all these error terms will remain largely constant with increasing PAN VMR, meaning that the percentage error will be lower than 43 % for profiles with mid-tropospheric values above 0.5 ppbv, but higher than 43 % for profiles with mid-tropospheric values less than 0.5 ppbv.

Index	Uncertainty	Nature	Est magnitude for profile with 0.5 ppbv PAN in the mid-troposphere	
			[ppbv]	[%]
1	Instrument noise	Random	0.15	30 %
2	Bias from a priori	Systematic	Depends on a priori	-
3	Absolute instrument calibration	Systematic	Assumed negligible	-
4	Spectroscopic Uncertainty	Systematic, direction unknown	0.04	8 %
5	H_2O	Pseudo-random	0.1	20 %
6	O_3	Pseudo-random	0.1	20 %
7	N_2O	Pseudo-random	0.05	10 %
8	Surface temperature	Pseudo-random	0.02	4 %
9	CFC-12	Systematic	+0.08 or less	+16 % or less
10	Emissivity (snow/ice)	Systematic	+0.1 or less	+20 % or less
11	Emissivity (silicate)	Systematic	-0.1 or less	-20 % or less
	Aggregate of 1-8		**0.22**	**43 %**

p. 5360, Section 4. I suggest major revisions to Section 4. The data presented in Figs. 8 and 9 are unconvincing because there is simply no apparent spatial pattern except perhaps a tendency towards high PAN concentrations at high latitudes (and whether or not this pattern reflects actual PAN concentrations is not at all clear). It is also problematic that for the particular TES observations (pixels) where elevated PAN is observed, there is apparently no spatial consistency, i.e. adjacent TES pixels show sharply different PAN concentrations. It would be much more convincing to present a case study based on a single TES overpass of a known biomass burning plume, such as the plume mentioned in Fig. 5. If the PAN retrieval algorithm has actual skill, one would expect that observations over a plume should at least demonstrate a reasonable pattern of low PAN levels outside the plume and high concentrations within the plume. Such a case study would not require in-situ measurements and should be feasible with data that were already processed.

We thank the reviewer for this helpful suggestion. We agree that showing examples of elevated PAN in known plumes is a convincing demonstration of retrieval skill. We have now added a Figure (Figure 9) that shows three examples of TES overpasses of known biomass burning plumes. This figure is reproduced below:

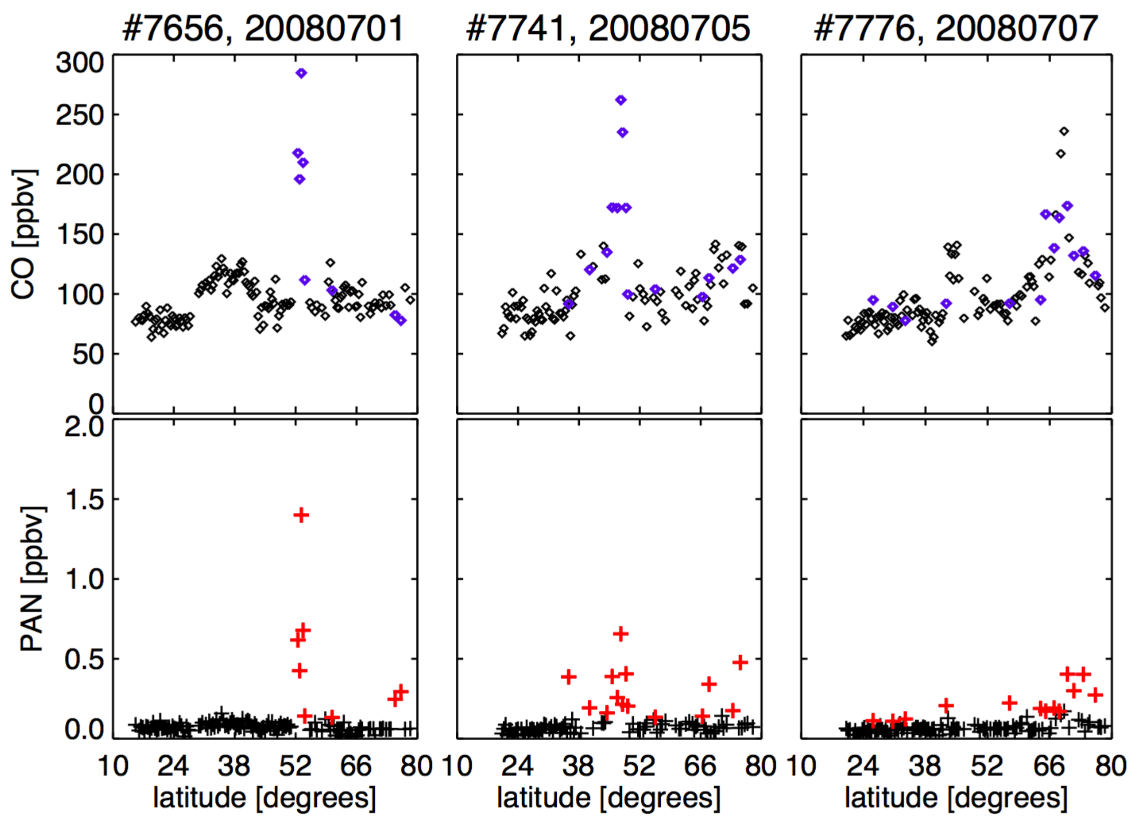

Figure 9: **Examples of elevated CO and PAN in boreal burning plumes (previously identified by Alvarado et al. (2010)) seen in TES special observations made during the July 2008 phase of the**

ARCTAS campaign. Colored points show the cases where the DOFS was greater than 0.6 for the PAN retrieval.

Technical Corrections

p. 5348, l. 11. 'Pacific' without 'Ocean' sounds colloquial

Updated.

p. 5350, l. 13. 'the the'p. 5354, l. 17. 'a prior' should be 'a priori'

Corrected.

p. 5356, l. 5. sentence including ' ... were used here for truth here ...' is awkward

This sentence has now been updated.

p. 5356, l. 25, p. 5359 l. 1, and p. 5360, l. 1 - misspelled 'uncertainty'

Corrected.

p. 5357, l. 19. misspelled 'evaluate'

Corrected.

p. 5358, l. 14. misspelled 'stability'

Corrected.

p. 5359, l. 6. capitalize 'rms' for consistency

Updated.

Anonymous Referee #2

Received and published: 17 July 2014

The article "Satellite observations of peroxyacetyl nitrate from the Aura Tropospheric Emission Spectrometer" by Payne et al. presents a novel look at the potential for tropospheric observations of PAN from the TES instrument. This work presents tropospheric PAN observations, an important addition to the study of the nitrogen cycle. The observations complement the established upper troposphere and lower stratosphere measurements of PAN made by limb-viewing instrumentation.

The paper presents a detailed look at an optimal estimation approach to invert TES measurements into PAN partial columns and makes a fairly convincing case based on simulations. Although, as the authors mention, in-situ measurements of PAN are sporadic, they do exist. I understand that the paper aims to convince the reader of the quality of the PAN data through simulations but I believe the paper could be improved if even a limited number of comparisons to aircraft measurements was made, for example the 2011 BORTAS campaign. Overall, the paper is fairly well presented and I'm happy for the paper to be published once the specific comments are addressed.

Specific comments

P5348 L5: It would be useful to mention which pressure/altitude range at which TES measurements are most sensitive to PAN. Even though you mention one degree of freedom, most information appears to come from the mid-troposphere looking at the averaging kernels.

Yes, most information does come from the mid-troposphere. The abstract will be updated accordingly.

P5351, L18: what are the advantages in using a natural log retrieval for the retrieval. Are you not falsely constraining the retrieval to have positive values rather than allowing negative values which, although physically impossible, are mathematically correct. Does using a log retrieval also mean that Gaussian statistics are still applicable?

Yes, we are falsely constraining the retrieval to have positive values, rather than allowing the retrieval to have negative values. This is a matter of choice. An argument often made in favour of allowing the retrieval to have negative values is that when retrievals are averaged, a zero result is made possible. In this case, TES can only detect values above ~0.2 ppbv. We are not able to say whether the retrieval is zero or very small, and we have included words of caution on the merits of averaging the TES PAN, since our "good" retrievals can never capture the low-PAN cases.

The retrieval in log space means that the Gaussian statistics apply in the log space. This leads to an inherent assumption about greater variability at the high end of the distribution of VMR values.

P5352, L6: The meaning of Sa is not clear. I understand Sa to mean the a priori covariance matrix. Could you please clarify.

Thanks for pointing this out. We should not have labeled this matrix as S_a^{-1}. It is the constraint matrix, but it is not an inverse prior covariance matrix.

In general, the constraint matrix in the retrieval can be different from the a priori covariance used in the calculation of the retrieval errors (e.g. Bowman et al., 2006).

For PAN, we do not have trustworthy a priori covariance estimates, at least certainly

not on a global scale. The error estimates presented in Table 2 were calculated empirically and did not involve the use of an assumed a priori covariance for PAN. At this stage, we do not have a trustworthy estimate of the PAN covariance. If it were known, we would use it. In theory, we could base such an estimate on GEOS-Chem results. In practice, we don't necessarily know how realistic this would be. Our experience has been that the models tend to underestimate variability. We might hope that as we continue to work with the PAN retrieval product, we can build up a better sense of the PAN variability.

P5352, L8: how do you quantify what is "relatively linear"?

Please see response to the same question from Reviewer 1.

P5352, L13: did the authors try to perform their own retrievals of P/T, H2O and the other important contributing gases within the PAN microwindows themselves? This could be done using the level 2 data as the initial state, meaning that the PAN retrieval isn't just fitting residual noise or another gas that hasn't been retrieved such as CFC-12.

The PAN microwindows are not the ideal regions to determine temperature, water vapor or the other contributing gases. We choose to take our information on these parameters from retrievals in other spectral regions that are better targeted towards those retrievals.

P5352, L20: Is it possible to verify that the surface emissivity varies linearly across the PAN spectral region? Is this not dependent on surface type? Have verifications been made against databases such as version 2 of ASTER (http://speclib.jpl.nasa.gov/) which compiles over 2400 spectra of natural and man-made materials.

We thank the reviewer for raising this important point. Reviewer 1 also raised this. Please see responses to Reviewer 1's comments. Note that the UW/CIMSS emissivity database referenced does draw on information from the ASTER spectral library.

P5353, L4: The authors state that the PAN signal is relatively weak compared to the noise. Have the authors attempted to co-add spectra over areas with similar emissivity, to reduce the noise on the spectrum? Or does the surface temperature and atmospheric variability make this unfeasible?

We did consider trying the approach of co-adding spectra. However, the atmospheric variability of PAN, combined with the spatial sampling characteristics of TES, make this tricky. TES "global survey" observations are spaced ~180 km apart along the orbit track. The TES "step and stare" special observations utilized in the new Figure 9 are more closely spaced, but even in that case, we see that the atmospheric PAN could vary greatly between one observation and the next. Co-adding spectra might be a possibility for instruments such as IASI, AIRS or CrIS where the spatial coverage is more comprehensive.

P5354, L9: Why was the initial guess profile set to a single value of 0.3 ppbv throughout the troposphere? The GEOS-Chem model results in Figures 3 and 4, for example, show a large variability in PAN with altitude. Wouldn't an average model profile, with associated variability provide a better starting constraint?

For the a priori constraint vectors, we did indeed use average GEOS-Chem model profiles (although not the associated variability, since we felt that the variability of PAN is not currently known well enough to say whether the model variability would provide a reasonable constraint). Provided the initial guess value is set high enough to allow jacobians that are non-zero, then the exact value of the initial guess should not affect the result of the retrieval. We acknowledge that a constant value of 0.3 ppbv throughout the troposphere is not a "realistic" PAN profile, but we set the initial guess this way in order to enable the retrieval to generate non-zero Jacobians at all retrieval levels within the troposphere. As stated in the manuscript, we choose to have an initial guess that is uniform with altitude in an attempt to avoid forcing the retrieval towards surface-maximum vs maximum-aloft profiles.

P5354, L12: Could the authors please explain "null-space" within the text.

We have added some wording to better explain what was meant here. Please see response to reviewer 1.

P5355, L12: The work of Allen et al. shows that PAN has many absorption bands across the IR range. Would extending the retrieval to these other bands improve information on the vertical structure?

As the reviewer points out, there are other PAN features in the thermal IR range. The other strong bands of PAN are located at around 800, 1300, 1730 and 1840 cm^{-1}. TES retrievals at 800 cm^{-1} are not possible due to higher instrument noise in the TES 650-900 cm-1 range. The PAN feature at 1300 cm^{-1} is strongly impacted by interference from methane and water vapor, which overwhelm the PAN signal. The 1730 and 1840 cm^{-1} PAN absorption bands are outside the spectral ranges measured by TES. These bands would be covered by, say, IASI, but would also be very strongly impacted by water vapor interference – much more so than the 1160 cm^{-1} band.

We have added a paragraph on these other bands to the end of Section 2.

P5356, L11: As biomass burning is a very large source of PAN, a useful application of the dataset would be to observe these events. Would TES PAN retrievals be possible within wildfire plumes (generally cloud-free, but high optical depth)? P5356, L16 makes a qualitative assessment of the capability of PAN retrieval in higher optical depths, but is not explored. Could the authors make a more quantitative assessment? Siberia would be a good test case, as an area with high PAN (based on the model) and a large number of wildfires at that time of year.

Yes, TES PAN retrievals are possible in wildfire plumes. Alvarado et al. (2011) had specifically looked at PAN detection in wildfire plumes originating from both Siberia and North America.

When the reviewer says "generally cloud-free, but high optical depth", we assume the reviewer is thinking about aerosol optical depth. In general, we do not expect aerosols to have a significant impact at this wavelength region. In particular, smoke aerosol particles are generally sub-micron in diameter and therefore small in comparison to the wavelength of the measurement. Scattering from smoke aerosol would not be expected to impact the TES radiances. To the extent that the aerosols could impact the retrieval, they would be spectrally smooth, would be interpreted as additional cloud optical depth and would be absorbed by the pre-PAN cloud/emissivity retrieval step.

P5357, L20: Are surface temperature and surface emissivity important components of the total error?

For a discussion of the surface emissivity error, please see response to Reviewer 1's comments on this topic.

The surface temperature is not a large component of the total error. We assessed the surface temperature error using the simulations (as was done for H2O) using a random perturbation with RMS 1 K. The surface temperature error should be around 0.02 ppbv. This estimate will be added to Table 2.

P5359, L23: I am slightly concerned by the assumption of CFC-12 being very well-mixed, particularly as Figure 1 shows that both CFC-12 and PAN have similar broad spectral shapes across a similar wavenumber range. Have the authors looked at the difference in the PAN retrieval with and without inclusion of CFC-12? Does changing the assumption of tropospheric CFC-12 vmr significantly affect the retrieved PAN vmr? What is the variability of the CFC-12 vmr shown across the NOAA Halocarbons & other Atmospheric Trace Species Group (HATS) surface sites (http://www.esrl.noaa.gov/gmd/hats/combined/CFC12.html) for April 2008?

Yes, we did look at the difference in the PAN retrieval with and without inclusion of CFC-12. The CFC-12 has a big spike in the middle that the PAN does not, making it straightforward to determine the presence of CFC-12 in the spectral residuals. CFC-12 is well-mixed and is seen in all observations. If PAN retrievals are attempted without CFC-12 in the forward model, then the retrievals will tend to indicate that elevated PAN is present in all observations.

The important point here is the low variability of tropospheric CFC-12 volume mixing ratios.

Long-term, globally-distributed surface measurements from the NOAA Earth Science Research Laboratory (ESRL) Global Monitoring Division (GMD) indicate that differences between measurements of CFC-12 in different parts of the globe are less than 3%. There has been a decrease of ~8 % in tropospheric values between 2005 and 2015. In the current TES climatology, we use a CFC-12 profile that is constant with time. TES climatology values, based on 2004 values, are on the high side.

We tried taking a 10 % reduction in CFC-12 as a pessimistic estimate for a sensitivity

test. If the true CFC-12 were indeed 10 % lower than the assumed CFC-12, this would lead to an overestimate of mid-tropospheric PAN by 0.08 ppbv. This error estimate has been added to the revised Table 2. (See responses to previous comments.)

This error could be reduced by implementing a time-varying CFC-12 climatology in the TES algorithm.

P5360, L10: This sentence is repetition of section 3.4.

This sentence has been removed.

P5360, L17: Would BORTAS aircraft measurements from summer 2011 provide suit-able validation data?

The BORTAS aircraft measurements are certainly an interesting dataset. TES did perform special "step and stare" observations during the BORTAS time period. The ideal validation measurements would be a set of profiles, covering the range of altitudes over which TES is sensitive, that capture examples of plumes and are extremely well co-located with the TES observations. That is a lot to ask for. Even with these special dense TES observations, spatio-temporal coincidence between TES overpasses and BORTAS flights was limited. This makes a true validation less than straightforward. We feel that TES/BORTAS comparisons merit their own dedicated study. The "ARCTAS plume" figure added in response to Reviewer 1's comments at least now demonstrates the ability of TES PAN to give feasible in-plume/out-of-plume results. For this Figure, we are able to draw on the previous work done by Alvarado et al. (2010) that identified fire plumes in TES observations.

We do have interest in further exploration of possible validation datasets, including BORTAS and the FRAPPE campaign that was conducted over Colorado in summer 2014. (Work is currently underway to process TES radiances for the FRAPPE time period, following a recent switch to a back-up laser on the TES instrument.)

We have updated the text to provide further information on what would be needed for a "true" validation, and have referenced the fact that TES did take special observations during BORTAS.

P5361, L21: The authors note a large number of retrievals over the Arctic, with some of the largest vmr. Can the authors please characterise any problems with the Arctic retrievals. In particular, does ice cause issues in terms of emissivity or characterising surface temperature? Does this impact your error estimate on the Arctic PAN data? Do you get the same sensitivity to PAN over the Arctic as for mid-latitudes (i.e. do the averaging kernels appear similar)?

In April, we expect the surface for almost all points above about 60 N to be snow or ice. (The GMAO surface temperatures are all below freezing.) We have not so far found particular problems with the Artic retrievals compared to other parts of the globe.

We looked at the chi-squared as a function of latitude and found no particular

dependence. The extreme high values shown on the map have chi-squared values that pass quality control.

We looked at the asymmetry of the retrieval residuals over the PAN spectral region (as an indicator of problems associated with surface emissivity – please see the new Figure 8) and found no behavior particular to high latitudes that would lead us to believe that the Arctic PAN retrievals are less trustworthy than other regions.

The averaging kernels do vary spatially. We would expect to see greater near-surface sensitivity for cases where the PAN profile peaks close to the surface. Also, for cases where PAN is high low in the atmosphere, the near-surface sensitivity tends to be greater in regions with higher surface temperature/greater thermal contrast. For cases where PAN is only elevated in the mid-troposphere, we would not expect near-surface sensitivity in any case. In the Arctic, we would generally not expect to see PAN close to the surface (unless we happened to observe right over a fire location, so we would not generally expect near-surface sensitivity.

We acknowledge that further investigation will be needed before drawing any strong conclusions about the high values of PAN observed in the Arctic. However, we do note that (a) high PAN values have been observed by aircraft in this region, as referenced in the manuscript and (b) PAN is more stable at colder temperatures, which means that at high latitudes, it does have a better chance of making it into the upper troposphere and sticking around.

Technical comments

Figures: Figure 1: Please change “F12” to CFC-12, as this is how the gas is referred to in the text. Also, N2O and NH3 are similar colours and so it is not easy to see which is most important within the figure.

Done.

Figure 5: The font size is very large compared to the text. Please reduce the font size.

Done.

Figure 8 & 9: The colours for East Asia do not appear to match, between figures. Please rectify this.

Done.

P5348 L2: does the algorithm have a name? This would be useful for future papers on TES PAN retrievals.

The retrieval approach presented here will be implemented in the TES V07 algorithm. We have added this information to the final paragraph of the conclusions.

P5356, L5: This sentence is jumbled and repeats "here", please rectify.

Rectified.

P5356, L25: Please change to "uncertainty"

Done.

P5359, L1: Again, please correct the spelling of "uncertainty"

Done.

P5360, L1: Please change to "uncertainties"

Done.