We would like to thank the referee for their helpful comments and suggestions. Included below is each of the referee's comments (italics) followed by our reply.

Responses to Referee 1 (Alexei Rozanov)

General Comments

Authors use outdated versions of OMPS Level 1 data (v2.0-2.4) although the new data version v2.5 is available already since May 2017. As version 2.5 already includes the pointing correction described in Sect. 3 of the manuscript this section would not be necessary any more if new Level 1 data was used.

Reply: It is true that a similar pointing correction has been included in the v2.5 L1G product, we have mentioned this in the revised manuscript. However we feel it is important to include the section as the manuscript serves as a description of our v1.0.2 retrieved ozone product which has already been used in several studies. We are planning on producing a new version of our ozone data product in the future based on the new L1G product but it is not feasible to include it here due to the computational burden of reprocessing the entire mission.

Pointing accuracy is mentioned as the main error source and the corrections in the order of 200-300 m seem to be considered by authors as important, oth- erwise one would rather skip Sect. 3. On the other hand, the authors do not hesitate to neglect the field of view of 1.5 km without making any considerations about the impact of this decision. As the field of view illumination is vertically inhomogeneous, I assume the neglect of field of view integration should have a similar effect as a misspointing. In this regard it is not quite clear why a very good agreement with MLS is still achieved and if the entire verification results might be accepted as trustable. To my opinion the evaluation must be repeated taking into account the field of view of the instrument.

Reply: We agree that neglecting the field of view has the potential to have an effect on the retrieval, however this is no trivial matter. The referee states that the field of view is 1.5 km, which is approximately true for a single pixel on the detector, but the provided level 1 product is gridded, bilinearly interpolated from four neighboring pixels. The actual field of view (both magnitude and shape) varies as a function of altitude and wavelength depending on where each pixel is, and this information is not publicly available. Furthermore, neglecting the instrumental field of view is a common assumption in many limb retrievals (including the operational NASA OMPS-LP retrieval and the OSIRIS retrieval, which this work builds upon) and we do not feel it is within the scope of this manuscript to perform a full study on this effect.

However, as stated in the manuscript we do intend to investigate this further in a future version of the retrieval. We expect that neglecting the field view has a vertical smoothing effect, which is why we do not trust our predicted 1 km vertical resolution and instead estimate it as 1–2 km. Our preliminary tests have indicated that neglecting the field of view has the potential to introduce a 1–2 % high bias near 20–25 km. This is considerably less than the $\sim$7% error you could see with a 300 m pointing shift at high altitudes.

As an improvement of the retrieval quality by using a 2D retrieval is a key topic of the manuscript, synthetic retrievals as done in Sect. 5.1 need to be presented for the whole orbit. This is necessary to assess if smoothing out the small latitudinal variations by 2D retrieval as seen around 50° S in Fig. 8 is a general drawback of this approach or just an insignificant outlier. Furthermore, a similar study should be performed for another season with a vortex edge in the northern hemisphere. This will allow the reader to assess how the viewing geometry affects the relative performance of the 1D and 2D retrievals. Another important question is how the retrieval results depend on the ozone distribution used to initialize the radiative transfer model. This question has not been addressed in the manuscript at all.

Reply: We have changed the requested figure to show the entire orbit, and also added a second orbit using synthetic data with northern hemisphere ozone depletion. Internal tests have shown that the dependence on the profile used to initialize model is negligible ($\leq 1\%$), we have added a statement to the revised manuscript to this effect.

The retrieval description is too much general with a lot of details hidden from the reader. For example, no or only insufficient quantitative information is provided about the latitudinal grid, reference tangent height and regularization parameters γ in Eq. (2) and α in Eq. (3)). The authors state that the a priori state vector is set to zero but make no comments about the values used to initialize the radiative transfer model. Are they also zero at the first iteration? The valid altitude range of the retrieval in not clearly identified.

Reply: We agree that the retrieval description was too general in many places. In the revised manuscript the description of aspects of the retrieval such as the state vector are given directly for the retrieval as applied to OMPS-LP, rather than first for a theoretical retrieval. Many of these specific changes are outlined as replies to the referees other comments below.

The validation is not sufficient to demonstrate the overall performance of the algorithm. The monthly mean comparison plots similar to Fig. 10 must be provided for absolute values rather than for anomalies for several latitude bands (tropics, middle and high latitudes).

Reply: The provided comparisons are intended to demonstrate the validity of the technique and not be a full validation of the dataset, which we feel is beyond the scope of this manuscript. As stated in the manuscript the validation work presented is preliminary, and a full validation is planned for a future paper.

Detailed comments

Page 2, line 24 *"... $\gamma_i I$ might be included..." - please make a clear statement if this term is included in your retrieval or not, if yes, what is the starting value and a typical end value of γ_i?*

Reply: We have added the statement that a small $\gamma_i = 0.1\cdot$ the mean value of the diagonal of $\mathbf{K}_i^T \mathbf{S}_\epsilon^{-1} \mathbf{K}_i$ term is included to primarily aid with the stability of the inversion.

Sect. 2.2 *State vector is described insufficiently. Both altitude and latitude grids must be specified exactly providing the upper and lower limits as well as the sampling.*

Reply: This section has been rewritten, we have opted to immediately explain the state vector for ozone here rather than a generic state vector. The horizontal grid is not spaced in latitude, rather it is in angle along the orbital plane of OMPS-LP, we hope this is now clear. Upper and lower limits and sampling have been noted for both the vertical and horizontal components in the revised section.

Page 3, line 13 *"A consequence of the limb viewing geometry..." - this is not a general consequence of the limb viewing geometry as a scanning instrument can be operated to avoid this problem (e.g. SCIAMACHY). This is rather a con- sequence of the imaging technique (2D detector array) used in OMPS.*

Reply: We have changed the wording to "A consequence of the OMPS-LP viewing geometry"

Page 3, paragraph starting at line 16 *this is an unnecessary general discussion which do not provide any useful information. It is highly questionable if the method described by authors is really that general as no references are provided. Furthermore, possible griding issues vary with the observation method. For example the issues are completely different if a combination of measurements along and across the flying direction is used. I recommend to remove the paragraph and focus on the detailed description of the setup used in the retrieval rather than discussing any "general" approaches.*

Reply: We have removed this information.

Page 4, Sect. 2.3, starting from line 16 till the end of the section *to my opinion this text does not provide any useful information as for the retrieval/modeling description it is absolutely irrelevant whether the model performs the internal transformation of the coordinates or not. If you think it is important you need to describe it in much more details to give the reader* *understanding what is performed, how and for what reason, and which implications it can cause. Otherwise the text must be deleted as in its current form it is just confusing.*

Reply: We have moved the figure and the text mentioning the mismatch between the instrument line of sight and the retrieval grid to the state vector section. We have also removed the text about the coordinate transformations as requested, and now simply state that the line of sight plane is projected onto the retrieval grid.

Page 5, lines 1-2 *"The sparsity of the Jacobian matrix can be improved..., as is done in Livesey et al. (2006)" - there are a lot of things which "can be done". The essential information is, however, if it "is done" in your retrieval or not. Please provide the numbers if it is done or clear statement that it is not done otherwise.*

Reply: We have added the statement "For our retrieval we limit each measurement to contribute to profiles within 10° of the tangent point.".

Page 5, lines 3-6 *This text does not contain any useful information. The matrices to be stored and inverted are already known from Eq. (2), their dimensions are already discussed in the first paragraph of the section, the fact if you solve the linear equation system using a solver for sparse or dense matrices is an absolutely minor technical information and a calculation of a memory space needed to store a 10000 $\times$ 10000 matrix is a very simple arithmetical exercise which is not relevant for a scientific paper.*

Reply: We agree that obviously calculating the storage requirements for a matrix is simple, but we do not agree that the information should be removed. One of the limiting factors for grid spacing, number of measurements used, etc., for a tomographic retrieval is computational. Other papers describing tomographic techniques such as Livesey et. al. 2006 and Ungermann et al. 2010 include similar types of information.

Ungermann, J., Kaufmann, M., Hoffmann, L., Preusse, P., Oelhaf, H., Friedl-Vallon, F., and Riese, M.: Towards a 3-D tomographic retrieval for the air-borne limb-imager GLORIA, Atmos. Meas. Tech., 3, 1647-1665, https://doi.org/10.5194/amt-3-1647-2010, 2010.

Livesey, N. J., Van Snyder, W., Read, W. G., and Wagner, P. A. (2006). Retrieval algorithms for the EOS Microwave limb sounder (MLS). IEEE transactions on geoscience and remote sensing, 44(5), 1144-1155.

Page 5, last paragraph *the paragraph is quite confusing. It not strictly defined what you understand as a "forward model run". In any case you have to simulate the radiance for every measured pixel, otherwise you just loose the information. Formally you can do just one "forward model run" and simulate everything. Thus, to understand this discussion, the reader has to know what is meant as a "run". Normally the forward model is run for* *each internal grid point, this might coincide with the location of the image or not. Surely a reduction of grid points reduces the computation time. So, actually, you just need to provide the information on the latitudinal grid and skip the remaining discussion.*

Reply: We realize that this section was confusing for those who are not familiar with SASKTRAN. The point we were trying to convey is that the expensive part of the radiative transfer calculation is calculating the multiple scatter source function, J_{MS}, which is a function of space, atmospheric state, and time. Once we have J_{MS} the final line integrals take little effort. The potential for computational time saving here is that since the grid points are close to each other, we can calculate J_{MS} in a spatial region that covers multiple grid points. The problem is that each "run" of SASKTRAN is for a single instant of time, so doing that is not strictly valid because each measurement obviously does not occur at the same instant in the time. So in this section we attempted to describe the potential issues with assuming that ~ 5 measurements occur at the same instant of time (roughly 100 s), it has nothing to do with changing the number of grid points (the largest issue is that the sun is assumed to be in the same location over this 100 s). We have rewritten the majority of this section to try to make this more clear.

Page 6, line 1 *"10° cone" - commonly the term "cone" is used for a 3D object while you have a 2D approach. Please use a proper notation. Furthermore, it is unclear how this "cone" is defined, I suppose from the Earths center, but it should be clearly stated to avoid a confusion.*

Reply: We have added the statement that the cone's vertex is the Earth's center, however in this case we believe that cone is the correct term. While the retrieval is a 2D approach, as soon as the atmosphere is allowed to vary in a second dimension (other than SZA) the symmetry in the source function is broken and it becomes 5 dimensional (position, direction) rather than 4 dimensional (altitude, SZA, direction). The source function is solved within this three dimensional cone.

Page 6, line 3 *"Each image..." - do you mean that the solar zenith angle changes from image to image? It is actually obvious that the illumination and composition of the atmosphere changes from one location to anther. Why is it an issue?*

Reply: We hope that this is clearer now that we have rewritten this section. The issue is that each measurement happens at a different time, a single SASKTRAN-HR calculation is one instant of time, so modelling multiple measurements with one SASKTRAN-HR calculation involves an assumption.

Page 6, line 4 *"... internal atmosphere is specified as a plane" - I suppose you mean the* *meridional direction. It should be clearly stated to avoid a misinterpretation.*

Reply: This has been changed to “as a plane in the along line of sight direction”.

Sect. 2.5 *Actually I did not find anywhere a statement about the variable defining the along-orbit grid, is it latitude, solar zenith angle, of anything else?*

Reply: The grid is the angle within the orbital plane, we hope that in the revised manuscript this is clear.

Sect. 2.5 *The last paragraph does not contain any useful information as it is not discussed how the OSIRIS images are compiled and how the corresponding radiative transfer calculations are done. Surely the listed conditions are not an issue for 1D retrievals if each observation is processed independently. I recommend to remove the paragraph.*

Reply: We think the modified section and our previous answers has made this clear. The issues are very much the same for the OSIRIS 1D retrieval which assumes that each scan happens at one instant of time, rather than running a new radiative transfer calculation for each individual observation.

Sect. 2.6 *Remove the first two paragraphs of the section. These paragraphs pretend to provide an overview of the methods fail however to do that as the discussion is to sketchy. Furthermore, this information is not needed for the discussion below.*

Reply: We have removed these paragraphs.

Page 6, line 28 *“For our retrieval ...”: please bear in mind that $\gamma_i I$ also works as a regularization term. So, when using Levenberg-Marquardt approach it is incorrect to state that the retrieval is completely unregularized. By the way, it is still not clearly stated if you use the Levenberg-Marquardt term in you approach or not.*

Reply: We do not believe this is correct in the standard use of the term “regularization”. The Levenberg-Marquardt term does not appear in the cost function as would a traditional regularization term, and in theory, the retrieval should converge to the same solution (neglecting issues of multiple local minima) with or without the Levenberg-Marquardt term. It is true that the Levenberg-Marquardt term can have a regularization effect if the retrieval is stopped before proper convergence, but that is not the case here.

Page 7, Eq. (3) *Provide α value.*

Reply: Added.

Page 7, Eq. (3) *The statement " **0** indicates a number of zeros equal to the number of altitude grid points" is wrong. It must be the number of altitude grid points minus one.*

Reply: Thank you, this has been corrected.

Page 7, line 4 *There are certainly some good reasons to use zero a priori state vector especially when employing smoothing constraints but the "simplicity" is not really the best one. It should be also mentioned that usage of zero a priori state vector often results in a low bias of the solution.*

Reply: We have removed the word "simplicity". We agree that with certain forms of regularization a zero a priori results in a low bias, however we have not seen anything to suggest that a second derivative constraint results in a consistent low bias. If there is a study that shows this we would be happy to state this and add a reference.

Page 7, lines 9-10 *I do not agree that the resolutions of the vertical and horizontal grids are strictly coupled. In principle any grid combinations can be used, this might require however a stronger regularization as the total amount of information remains the same. The main challenge here is to identify the optimal set of grids and regularization parameters. This set might however depend on the targeted usage of the retrieval data.*

Reply: I think we are mostly in agreement here, when we say the resolutions are coupled we meant to refer to the resolution of the retrieval, not the actual grid. We have changed this to state "the retrieval vertical and horizontal resolutions are inherently coupled together". The main idea we meant to convey is that crudely if we reduce the retrieval vertical resolution, there is more information available for the horizontal part.

Page 7, lines 12-13 *"The effect of a one dimensional retrieval on horizontal regularization...."*

- *I guess you mean "horizontal resolution".*

Reply: Thank you, this has been changed.

Table 1 *Please provide the reference tangent height for each interval.*

Reply: We have added the normalization altitude for each triplet to the table.

Sect. 2.7.1 *What is the minimum retrieval altitude for ozone?*

Reply: This information is now available much earlier in the revised manuscript.

Page 8, line 7 *Here and further below in the text you are talking about the "atmospheric upwelling". I suppose you mean the upwelling radiation. However, this notation is commonly used in the scientific community to describe the dynamic processes and means the upward moving air masses rather than radiance. Please use another notation throughout the text to avoid a confusion.*

Reply: We have changed all occurrences of "atmospheric upwelling" to "upwelling radiation".

Page 8, lines 20-21 *I guess Eq. (4) is valid for both triplets and doublets. "... for triplet k " in line 21 should be "... for triplet l ".*

Reply: Corrected.

Page 9 *"...any errors in the absolute calibration ..." - this is not completely true for an imaging instrument because the information for different tangent heights comes from different areas of the CCD and can have different calibration errors.*

Reply: This is true but we do not think our statement "helps to minimize any errors in the absolute calibration" contradicts this. The full line in the revised manuscript now reads "The high altitude normalization helps to minimize errors in the absolute calibration of the instrument and reduces the sensitivity to upwelling radiation."

Page 10, Eq, (6) *It is not clear how the second term is employed in the retrieval as the modeled Rayleigh background needs to be subtracted in the same way from both measured* *and modeled radiances and thus is canceled out when calculating* $\mathbf{y} - F(\mathbf{x})$ *in accordance with Eq. (2).*

Reply: Thank you, this was confusing in the text. The Rayleigh subtraction is not used in the actual retrieval, although as you pointing out it would have no effect, it is only used to determine the high altitude normalization location based on the procedure of Bourassa et. al. 2012. We have modified the text to make this clear.

Sect. 2.7.2 *No information is provided about how the aerosol extinction coefficient is calculated for other wavelengths.*

Reply: This is done using the same Mie code and assumed particle size distribution as for the phase function, the revised manuscript notes this and adds a reference to the source of the index of refraction data.

Page 11, line 1 *"... albedo is handled in a two-dimensional sense ..." - what is the second dimension for the albedo?*

Reply: We have reworded the first part of this sentence to "While albedo in the forward model is allowed to vary in the horizontal direction".

Sect. 2.7.3 *40 km tangent height to retrieve the surface albedo is quite high. Have you checked a possible influence of the stray light at this tangent height?*

Reply: We have done some internal tests here and have not noticed any significant effect on the retrieved ozone by changing the albedo retrieval height. Many past studies such as Loughman et al. 2005 have found that the absolute value of the retrieved albedo does not have a large effect on the ozone retrieval. Jaross et al. 2014 estimates the stray light contribution to be only 5% at 65 km for 750 nm, so we do not expect a large problem with using 40 km.

Sect. 2.7.3 *The influence of the albedo spectral dependence must be discussed. For example, for a green vegetation the albedo obtained at 745 nm can be very different from that at 602 nm (red edge).*

Reply: We have added a reference to Loughman et. al. 2005 which discusses possible errors associated with neglecting the spectral albedo dependence. However it is important to remember that the albedo is not really surface reflection and is merely an approximation for the unknown diffuse upwelling radiation, thus you would not expect as harsh of a spectral dependence that you would see with vegetation.

Sect. 3 *The section is unnecessary as all discussed corrections are already implemented in the Level 1 v2.5 dataset of NASA.*

Reply: See our reply to the general comment above.

Sect. 4 *If Levenberg-Marquardt term is used in the retrieval it must be also included in the error analysis.*

Reply: The Levenberg-Marquardt term does not appear in the cost function, and the error analysis is a linearization applied to the cost function so we do not see why this term should appear.

Sect. 4 *Is the signal to noise of 100 is used only in the error analysis or in the standard retrieval as well? Why was not the signal to noise data provided in Level 1 data set used? The latter would provide a realistic instead of maximum error estimation.*

Reply: The SNR of 100 is used in both. The OMPS L1G documentation states that the SNR provided is an "Estimate of detector noise and not an estimate of random measurement uncertainty.", and we were told by the NASA OMPS-LP team that a value of 100 is more realistic.

Page 13, line 9 *Only in the error analysis section the reader learn that the logarithm of the number density is the retrieval parameter rather than the number density itself. This must have been mentioned already in Sect. 2.2.*

Reply: The revised manuscript should correct this.

Page 13, line 14 *what does "but near where the tropopause lowers at midlatitudes" refer to?*

Reply: Reworded to "near where the lower bound of the retrieval changes (due to the lowering tropopause) at mid-latitudes."

Fig. 7 *Suboptimal color scale. How is the sign of the distance from the retrieval location defined?*

Reply: Color scale has been changed. The sign is negative towards the start the start of the orbit in the southern hemisphere, we have added this to the figure caption.

Page 14, line 3 *"Since the regularization term..." - once again, do not exclude the Levenberg-Marquardt term from the discussion.*

Reply: The Levenberg-Marquardt term does not have any effect here.

Fig. 7 *The definition of the vertically/horizontally integrated averaging kernels is not quite clear. You have a set of averaging kernels for each vertical/horizontal grid point and each of them spans in both vertical and horizontal directions. Is the integration done over these directions? Is yes you seem to show one averaging kernel at each altitude in each panel in Fig. 7? If it was true I would expect the plural in the beginning of line 8 as you show multiple averaging kernels for different altitudes in each panel of Fig. 7. If my understanding of the definition is correct, I would like you to explain why there is a clear maximum at 40 km in tropics and 45 km at mid-latitudes and how it can be interpreted in terms of the retrieval sensitivity.*

Reply: As for the plural vs not plural, Technically there is only one averaging kernel for the entire orbit and what is being shown are multiple rows. We have modified the text to make this clear.

The difference in peak altitude is a little curious, we have two possible explanations. The first reason is that lines of sight below the tropopause are not used in the retrieval. At mid-latitudes we have lines of sight going from 10–18 km, which for the strongest absorbing UV triplets have peak sensitivity in the 40–50 km region owing to the optically thick line of sight path. Since these lines of sight are missing in the tropics the sensitivity peak is lower. Another way of thinking about this is that generally the information content is poorest at the retrieval boundaries, and increases away from them. Since the lower boundary is higher in altitude in the tropics it makes sense that the information maximum shifts downward.

The second cause is the difference in solar zenith angle between the two points. For the OMPS-LP geometry, the tropics have low solar zenith angles with minimal solar attenuation compared to the limb path. Higher latitudes have higher solar zenith angles where solar attenuation becomes more important. It is expected that sensitivity overall shifts upwards in areas with significant solar attenuation as the attenuation happens above the tangent point.

Page 14, line 9 *"Only minor differences ..." - to my opinion the majority of differ- ences occur around 40 km and they are not minor.*

Reply: We see the referees point, however at 40 km the difference in FWHM is less than

25 km. We have reworded the text to state “Only minor differences in the FWHM ...”

Page 14, lines 11-12 *“it was found that ...” - it is hard to believe as it is widely known that the averaging kernels for “relative” retrievals (i.e. retrieval of relative deviations from a priori or logarithms) depend on the atmospheric state. Please provide the averaging kernel plot for different season to justify you statement.*

Reply: We have added a second orbit (from a different season) averaging kernel to the figure. We have also reworded the offending sentence to “it was found that deviations from orbit to orbit are small enough that the above resolution estimates are representative for the entire dataset.”

Fig. 7 *why does the tropics plot have a white area below 18 km, how is the lower boundary of the retrieval defined?*

Reply: The data is masked below the lowest retrieval point which is the first altitude above the tropopause, we hope this is clear in the revised manuscript.

Page 14, last paragraph *It is absolutely inappropriate to neglect the instrument field of view without any investigations as it might lead to a significant change in both the retrieval results and error analysis.*

Reply: See the reply to the general comment.

Sect. 5.1 *The results must be provided over the whole orbit as it is essential to estimate how the retrievals compare outside the vortex edge region. Another simulation for a different season with a vortex edge in the northern hemisphere needs to be provided to assess the influence of the viewing geometry.*

Reply: We have modified the first figure to show the entire orbit. We have also added a second simulation for a different season with a strong ozone gradient in the northern hemisphere.

Page 15, line 25 *“For limb scatter measurements ...” - please illustrate this by plotting the averaging kernel for about 65° S and 15.5 km in both horizontal and vertical directions using a proper color scale.*

Reply: While the averaging kernel does also somewhat show this effect (in fact, it only does because some regularization is present) we do not think the requested figure is appropriate to justify the statement “For limb scatter measurements ozone sensitivity is larger on the instrument side of the line of sight”. The averaging kernel is specific to our 2D retrieval, and we are talking about an effect that is retrieval independent. The proper figure is one of the two-dimensional Jacobian for a single line of sight, of which there are many in paper referenced in the text (Zawada et. al. 2017).

Sect. 5.2 *this section is not really informative and can be skipped. Details on the execution time suit better in the algorithm description section.*

Reply: We have removed this section and moved the information into the algorithm description section.

Sect. 5.3 *Not only the anomalies but also the monthly mean values themselves need to be compared. This needs to be done for different latitude bands (tropics, mid-latitudes, high latitudes).*

Reply: See the reply to the general comment above.

Fig. 10 *Why the altitudes above 59 km are not shown? If I understand it correctly, the retrieval runs up to 59 km.*

Reply: We assume the referee means why are altitudes above 50 km not shown. 50 km tends to be a common cutoff for ozone anomaly figures (and trend figures) due to the strong diurnal effect. In the coincident comparisons we can go above 50 km as the time difference between the two measurements is quite small. There are also issues where filtering MLS data according to the recommended procedure frequently cuts the data off ∼55 km or occasionally lower, so below 50 km the sampling is roughly consistent for both instruments.

Page 19, lines 1-3 *“... with the horizontal along-track resolution being poorer..” - please provide the values of the resolution and sampling for both instruments.*

Reply: We have added this information to the text.

Page 19, lines 5-6 *“.. has been degraded to the MLS pressure grid with a least square fit...” - please clarify what exactly was fitted and how you can degrade the vertical resolution using a least square fit. Here, a convolution with averaging kernels would be more suitable.*

Reply: This is simply the recommended procedure in the MLS data quality document, we have added a reference to the data quality document to indicate this. The OMPS-LP measurements are converted to pressure at native resolution in pressure, then rather than interpolating these values to the MLS pressure grid a least squares fit is done assuming linear VMR variations. We have opted not to apply the MLS averaging kernels (in addition to the least squares fit) since our vertical resolution (estimated 1–2 km) is not significantly better than the MLS vertical resolution (∼3 km). Futhermore, the MLS averaging kernels are fairly strongly peaked (peak values of 0.6 in the UTLS) so it would not be expected to make any significant differences. Jiang et al. 2007 did compare both of these methods (least squares fit vs averaging kernel) and found negligible differences even when comparing high resolution sonde measurements to a version of the MLS data with poorer vertical resolution than what we are using here.

Jiang, Y. B., et al. (2007), Validation of Aura Microwave Limb Sounder Ozone by ozonesonde and lidar measurements, J. Geophys. Res., 112, D24S34, doi:10.1029/2007JD008776.

Fig. 12 *Provide the lower and upper altitude of the plotted range. Provide the same plot from 1D retrieval. Explain the lower limit of the retrieval.*

Reply: Tick labels have been added for the maximum and minimum of the plotted altitudes. The lower limit is now explained earlier in the revised manuscript. We do not see any value in adding results from the 1D retrieval here. Our only claim made about the 1D retrieval is that it has problems in the presence of large horizontal gradients, of which there are none in this orbit.

It would be also interesting to show some examples from NASA Level 2 data, especially in Fig. 14.

Reply: We agree this would be an interesting study, and in fact we believe there is work being done by other groups on comparing OMPS-LP retrievals by different processors, but we feel it is beyond the scope of this manuscript to include these comparisons.

Technical corrections

Page 2, Eq. (1) *matrices have to be shown in bold face to match the corresponding notations in the text.*

Reply: Corrected, thank you.

Page 15, line 5 *duplicated word "those".*

Reply: Corrected, thank you.