

Interactive comment on “Ground based lidar and microwave radiometry synergy for high vertically resolved thermodynamic profiling” by M. Barrera-Verdejo et al.

Anonymous Referee #1

Received and published: 1 July 2015

An optimal estimation (OE) technique is used to merge Raman lidar returns and microwave radiances to produce profiles of absolute (or relative) humidity and temperature. The lidar data is processed externally while the microwave radiances are forward modelled. A priori information is drawn from local radiosonde observations. The performance of the algorithm is demonstrated with data from the HOPE field campaign, exploring the influence of the two sources both separately and together. The combination is shown to reduce the discrepancy of results against coincident radiosonde launches.

I cannot recommend the publication of this paper as it does not represent a relevant

C1793

or novel advancement of the field. Improving the vertical resolution of water vapour profiles is a worthwhile aim but my reading of Figs. 2, 5, and 6 is that this algorithm is in essence a complicated weighted mean of the input parameters. The degrees of freedom in Tab. 2 clearly show that the output is dominated by the lidar return, which is not calculated by this work. I would have expected there to be some merit to the product near the boundaries of the lidar-data regime due to the action of the a priori covariance, but the averaging kernels shown in Figs. 1 and 2 of the Appendix of the response to the Editor indicate otherwise (though such improvements would not be worthwhile on their own).

The ad hoc means by which the lidar data is introduced using a qualitative cut-off is a poor application of the OE technique. In my opinion the key advantage of OE is that the mathematics determines which data is relevant and which isn't based on the error covariance matrices independent of the investigator's preferences. The entire lidar profile should be input to the algorithm and the error covariance matrix properly assessed (e.g. quantifying overlap uncertainties). If computational cost is a concern, the profile can be truncated, but this should be done consistently across profiles and be higher in the atmosphere. Output products can be truncated based on their uncertainty (i.e. rather than cut off input data when the uncertainty is excessive, cut off the output when the uncertainty is excessive).

To be suitable for publication, I would expect the addition of a forward model of the lidar return (preferably ingesting raw profiles and taking other corrections as parameters but ingesting the attenuated backscatter coefficient would be acceptable, though the estimation of the error covariance would be more complex). I am familiar with Delanoe and Hogan (2008); Pounder et al. (2012); Povey et al. (2014); Sica and Haeefele (2015) having successfully constructed and used such models. Though none of these directly considered water vapour, the complexity of the lidar equations is negligible compared to the spectrographic models and radiative transfer equations on which optimal estimation is routinely applied. This would be a substantial task, but would represent a significant

C1794

advance for the community.

In the authors' favour, the piece was clearly written and easily understood. I especially appreciated the lengthy descriptions of each figure, which were well-chosen and clear (I agree with their sentiment that averaging kernels are best placed in supplementary material as, though vital to appreciate the function of an OE retrieval, they are confusing to inexperienced readers). The application to both an extended data set and individual observations is appreciated. I would encourage they extend their analysis to improve upon existing lidar water vapour products and submit a new paper.

Some comments that they may wish to consider during that work follow:

Throughout: You frequently use 'error' when you mean 'uncertainty'. Error refers to the difference between the value measured and the theoretical true value of the thing you are measuring. The uncertainty describes the distribution of that error as a means of quantifying the range of values that are consistent with the value you report. For example, discrepancies relative to the radiosonde are errors but the values from Eq. (4) are uncertainties.

5468, 24: These are excessively precise numbers considering they derive from the analysis of a single profile. One significant figure would more accurately represent your understanding of the change in error.

Sec. 1: The referencing in this section is poor, especially for a student that is preparing a thesis, as they primarily draw from the previous work of the authors. Some reference to previous applications of optimal estimation to lidar and/or microwave radiometry should be included, in addition to the specifics of Raman lidar analysis. Some other papers from the HOPE project would also provide useful context.

5469, 12: I thought we could resolve boundary layer turbulence with sodar and that the issue with cloud processes is the opacity of clouds, rather than the resolution of the observations?

C1795

5472, 19: Neither of these papers describe how the relative humidity is calculated. Considering the importance of the error covariance in OE retrievals (and your later conclusion that the uncertainty is underestimated), this technique should be briefly outlined alongside the procedure by which the uncertainty is estimated.

5470, 3: From Fig. 9, I can see why you would blame the overlap function on failures in the lidar (as they are concentrated near the surface), but considering water-vapour measurements generally derive from a ratio of measurements, you will need to present more evidence for this (e.g. some manner of sensitivity study) to convince most readers.

5472, 9: Raman scattering is stimulated by all three wavelengths. We generally observe 355 nm because the cross-section is greatest there, but the other wavelengths exist. For example Althausen et al. (2000) observes water vapour at 660 nm.

5473, 7: For those unfamiliar with lidar, it may be useful to point out that the different ranges result from the relative strengths of the signals used.

Sec. 3.1: It is unusual to use the Gauss-Newton iteration rather than the Levenberg-Marquadt. The model you use is straightforward, but the additional computational expense of the L-M method is generally minor considering the improvement in convergence (G-N being more susceptible to local minima of the cost function). Regardless, the convergence criterion is poorly described in this section. You should be consistent and use $F(x_i)$ rather than y_i .

Sec. 3.1: I initially thought you were using inefficient expressions, but later realised you are retrieving more states than data. It is advisable that the number of state vector elements n is less than the number of measurements m , as otherwise the solutions are underconstrained and can be sensitive to the initial conditions (or other retrieval parameters). I would be inclined to slightly reduce the resolution

C1796

of the profile through the lidar region (or use a consistent grid throughout) such that $n < m$. The loss of resolution in the lidar region may be worth the increased reliability of solution.

5475, 20: Influential though Rodgers (2000) is, you don't need to reference him three times in as many lines.

5475, 24: BT would be a vastly more common acronym for brightness temperature.

Eq. (6): I thought that it was standard to place the diagonal elements of a matrix from top-left to bottom-right? You are trying to be consistent with Fig. 1, but you should check the style for this journal (and invert Fig. 1 if necessary).

5481, 10: This paragraph, and many that follow discussing the lidar-only retrieval, are problematic. Since you have no forward model for the lidar data, the 'retrieval' without MWR data will be a weighted averaged of the profile and the a priori. That isn't a sensible product as the a priori uncertainty is large compared to the uncertainty on the lidar product. Either keep the current discussion but compare to the unprocessed lidar product, or concentrate more on the response of the retrieval where there is no lidar data (much as I would never trust retrieved values in those regions as there are zero degrees of freedom, the algorithm at least did something there).

5481, 27: Again, this isn't a meaningful statement. I would expect the error of the lidar data to closely follow the measurement covariance. It will be slightly reduced due to the a priori and the MWR data, but that isn't interesting.

5482, 3: I think statements such as this best demonstrate my problem with the algorithm's output. The results when using both inputs are only a 'best fit' by eye. There are substantially more data points in the lidar region than above, such that the improved fit in the free troposphere is mathematically irrelevant. In the lidar

C1797

region, the algorithm basically doesn't change the input data (as the uncertainty on it is small relative to the microwave and the a priori). I can understand that, qualitatively, you were aiming to use the a priori covariance matrix to propagate the lidar information vertically, but I do not feel the algorithm you present is a sensible means of doing that (nor does the text communicate an emphasis on improvement above the lidar profile, if that is the case).

Fig. 4: Though this is a nice plot, how does the lidar product before your processing compare? Also, such plots are always more convincing when paired with a scatterplot.

Sec. 4.2.2: To what extent are these radiosonde measurements independent? You imply they are used in the calibration of the lidar product.

Fig. 7: You say the uncertainty is increased by a factor of four but the retrieved uncertainty on the product only increases by a factor of two. Though it doesn't need to be mentioned in the manuscript, it would be useful to confirm that you multiplied the covariance matrix by 16 (or multiplied the uncertainty by four before squaring).

5487, 16: I have no problem with applying an empirical factor-four correction as a theoretical exercise, but there would need to be a more thorough error budgeting of the lidar product if you were to use this in practice.

5488, 7: Considering you retrieve order 100 values in the lidar region, 12-25 degrees of freedom aren't 'large'. They're simply more numerous than those from the MWR. In fact, I'd take the relatively low number of degrees of freedom as a sign that you should retrieve on a coarser grid in the lidar region as each data point there is not conveying much independent information. The choice of resolution could be guided by the width of the averaging kernels.

C1798

Fig. 8: I know this plot was inserted in response to the Editor's comments, but it's difficult to interpret degrees of freedom without knowing the number of elements in both the measurement and state vectors (which are never mentioned). I'd personally like to see the variation in width of the averaging kernels, but that can be in the supplementary material as a subplot to the kernels themselves.

5490, 16: It improves, but I wouldn't say 'strongly'.

5491, 26: The two profiles shown are consistent with each other, such that I'm not certain the 'improvement' relative to the radiosonde isn't coincidental. You would need to demonstrate improvement over at least a few test cases.

Additionally, a few language corrections I noted while reading:

5467: In the title, I believe 'high vertical resolution' would be better phrasing than 'high vertically resolved'. In fact, throughout the paper you use 'highly resolved' when I believe 'high resolution' is more appropriate.

5470, 29: sensors compensateses for

5471, 21: of the clouds'' lifetime

5473, 1: due to an insufficient overlap

5474, 27: of a horizontally

5475, 14: An Optimal Estimation Method allows one to

5481, 4: 'it allows one to work' or 'it can work'

5482, 12: At ground level, the only two available

5486, 19 As argued in Sect. 4.2.2

C1799

5489, 10: the OEM allows one to work

5490, 15: This sentence is repeated half-way through.

5492, 25: The synergy presents its

5493, 6: could be desirable.

Fig. 1: (from top to bottom and left to right)

Fig. 10: enclose the area where lidar

References

- Althausen, D., Müller, D., Ansmann, A., Wandinger, U., Hube, H., Clauer, E., and Zörner, S.: Scanning 6-wavelength 11-channel aerosol lidar, *Journal of Atmospheric and Oceanic Technology*, 17, 1469–1482, doi:10.1175/1520-0426(2000)017<1469:swcal>2.0.co;2, 2000.
- Delanoe, J. and Hogan, R. J.: A variational scheme for retrieving ice cloud properties from combined radar, lidar, and infrared radiometer, *Journal of Geophysical Research — Atmospheres*, 113, D07 204, doi:10.1029/2007jd009000, 2008.
- Pounder, N. L., Hogan, R. J., Várnai, T., Battaglia, A., and Cahalan, R. F.: A Variational Method to Retrieve the Extinction Profile in Liquid Clouds Using Multiple-Field-of-View Lidar, *Journal of Applied Meteorology and Climatology*, 51, 350–365, doi:10.1175/jamc-d-10-05007.1, 2012.
- Povey, A. C., Grainger, R. G., Peters, D. M., and Agnew, J. L.: Retrieval of aerosol backscatter, extinction, and lidar ratio from Raman lidar with optimal estimation, *Atmospheric Measurement Techniques*, 7, 757–776, doi:10.5194/amt-7-757-2014, 2014.
- Sica, R. J. and Haefele, A.: Retrieval of temperature from a multiple-channel Rayleigh-scatter lidar using an optimal estimation method, *Applied Optics*, 54, 1872–1889, doi:10.1364/AO.54.001872, 2015.