Atmospheric
Measurement
Techniques

Discussions

Interactive comment on "A practical method to remove a priori information from lidar optimal estimation method retrievals" *by* Ali Jalali et al.

Anonymous Referee #1

This paper applies the technique of von Clarmann and Grabowski (2007) to the optimal estimation of lidar profiles. An existing optimal estimation retrieval is performed at the native resolution of the lidar data. The averaging kernels from that retrieval are used to devise an optimal retrieval grid, whereby each retrieval bin contains one degree of freedom. Deviating from the original technique, a second retrieval is performed on that grid, using an arbitrarily large a priori uncertainty to minimise it's influence on the retrieval. The technique is demonstrated for three profiles, showing an increased range of useful results and a qualitative improvement in the agreement with an independent radiosonde.

This research demonstrates a marginal improvement in the vertical extent of the lidar profiles. Though I disagree with the author's insistence that a priori information is best

avoided, I recommend the paper for publication after minor corrections to clarify a few parts of their argument. The English is of a high quality throughout and I include some grammatical corrections and suggestions in an attached PDF.

In the following, PxLy denotes line number y of page x the submitted paper. My primary critiques are:

- Optimal estimation without a priori information is arguably not optimal estimation. The use of a prior for regularisation is what distinguishes the method from other manners of regularisation, such as smoothness constraints (e.g. Pounder et al. 2012, 10.1175/jamc-d-10-05007.1). I appreciate that the influence of the prior on the final solution is a focus of much critique of the OEM, so the removal of the prior is 'desirable' in one sense but it's an unfair representation of the information available. Why not change the grid but retain the prior, which should then have a minimal impact on the points you currently retain?

- As you point out on P2L28, it would have been preferable for you to use a different prior. I am not surprised that a single, globally representative profile is a biased prior. Other sources of information are available so why not use them? The reanalyses and forecasts of GFS and ECMWF seem like excellent candidates. If you wish a more climatological prior, there are statistical analyses generating a set of representative atmospheres, such as §3 of https://www.ecmwf.int/sites/default/files/elibrary/2008/11040-generation-rttov-regression-coefficients-iasi-and-airs-using-new-profile-training-set-and-new.pdf (which is, admittedly, a tad obscure). Or use your record of radiosonde launches to make something more locally representative. A globally representative profile makes sense for a moving lidar, but for the permanent installations you discuss, a more specific prior seems appropriate.

- In my opinion, the primary use of removing the prior from optimal retrievals is for averaging multiple retrievals together. Combining retrievals that contain a prior
gives that assumption undue influence on the average and biases Level 3 data. Had you considered this?

- Though you never clearly state so, I can see that adapting the retrieval grid to the information content provides a better representation of what the measurement actually told us. The paper would benefit from an explanation of why you prefer this statistical regridding approach to defining a single inhomogeneous, vertical grid that has decreasing resolution with height. I suspect it's because of dry layers, such as Fig. 5, that 'waste' state vector elements on heights where there is minimal information content available, but explicitly explaining that somewhere before §6 would be helpful. It would also be worth considering how using uneven grids makes it difficult to average retrievals and can confuse inexperienced users.

- In Figs. 5 and 9 please explain the meaning of the % unit. Is this fractional uncertainty (e.g. uncertainty ÷ value), fraction of the uncertainty (e.g. statistical uncertainty ÷ total uncertainty), or a conversion of the unit of mixing ratio? The presence of a point $> 100\,\%$ implies the first option, but it then makes little sense why the uncertainty tends to zero with height, as the magnitude of uncertainty should tend towards the value given in the a priori covariance matrix. I agree with your arguments at the bottom of P9, but if your prior is controlling the retrieval and uncertainty is tending to zero, I suspect that your a priori covariance is far too small at these heights.

- Throughout the experiments, you argue that the highest retrievals aren't viable. Could you please clearly state the conditions whereby you define a retrieval to be invalid? It appears to be based on the magnitude of uncertainty but it's unclear (a) if you are considering total uncertainty or the statistical component thereof, (b) where your threshold lies, as values from 60–100 % are mentioned, and (c) if the % there means the unit of mixing ratio or a fractional uncertainty.

- Figs. 6 and 10 are a somewhat unfair comparison. Considering the full resolution,

does the coarse grid retrieval have a lower RMS summed over the entire valid profile? Also, how do you interpolate the radiosonde profile onto the grid of each lidar retrieval? In theory, a fair comparison would use the averaging kernels of each retrieval to make a weighted average of the radiosonde for each profile and compare to that (e.g. Rogers and Connor 2003, 10.1029/2002JD002299).

P9L8 The red line first drops below 0.9 at about 2.2 km. I think what you meant to say is the red line *last* drops below 0.9 at 2.7 km.

P19L24 I strongly disagree with your implication that retrievals using a prior are invalid. The erroneous belief that prior information overwhelms statistical retrievals is too widely held and doesn't need assistance. Your retrievals were likely 'invalid' because of an overly general prior (possibly with unreasonably small uncertainties) on data too noisy to fully constrain the problem at that point.

P19L26 I also disagree that your paper shows that this removal method 'improves the validity of the retrieval.' You showed one example where the agreement qualitatively improves. On the one hand, you could have potentially achieved the same result through the use of a better prior. On the other hand, it may be that all possible priors are wrong, in which case you need to add another section to this paper.

- It would be preferable to make a statistical analysis of the change in RMS vs. radiosonde across many days of observations, as it is possible that beneficial results shown in the three examples are coincidental. No one ever provides such an analysis during revisions, but that won't stop me from asking.

Some more minor comments,

P2L14 It took a few minutes to work out that this paragraph was referencing two different figures, one of which isn't in this article. Since this article won't be published until

2019, it may be easier for future readers to distinguish, but it may just be easier to reproduce that figure here.

P2L22 The difference may be smaller than random uncertainties, but is it smaller than the a priori uncertainty?

Eq.6 You omitted the prior term. The equation should read,

$$\hat{\mathbf{x}} = \mathbf{x}_a + \mathbf{A}(\mathbf{x} - \mathbf{x}_a) + \mathbf{G}\epsilon = (\mathbf{I} - \mathbf{A})\mathbf{x}_a + \mathbf{A}\mathbf{x} + \mathbf{G}\epsilon.$$

The second form follows Rodgers (3.12), which I find a better illustration of the impact of the averaging kernel. If the averaging kernel is a unit matrix, the first term is clearly zero so the prior has no influence. However, your mileage may vary.

P8L2 It isn't 'necessary' to remove the regularization term; it's merely desirable. Perhaps word this sentence, 'We then remove the regularization term in Eq. 4 by choosing an arbitrarily large a priori uncertainty.'

P9L3 I can see why you want Fig. 5 after 3 and 4. Perhaps remove the reference here so the reader doesn't worry they've missed something from the out-of-order mentions?

P9L20 It may be better to say that these retrievals aren't 'useful' or 'meaningful', as 'viable' implies that something about the retrieval process failed. The retrieval still works, it's just that your observations are swamped by noise such that the retrieval isn't telling you anything useful.

P9L22 Why keep telling us only the statistical component of the uncertainty? The other components are important too!

Fig.4 The caption implies that the reason the points are rejected is their large span. The text implies they are rejected because of their large uncertainty. Please clarify.

Fig.6 The fine and coarse grid appear to have the same vertical coordinates (and number of points, which must be wrong).

P11L12 You mention a 'traditional method' which is not shown in Fig. 9.

P12L12 You didn't show invalid retrievals in Fig. 6 so why show them here?

Fig.13 Are you sure 'systematic' is the right term here? Optimal estimation models all uncertainties as Gaussian, which would be referred to as random errors. I know the various retrieval parameters are systematic in that they affect all points in the profile equally, but the definition of a systematic error is more like a bias and something that optimal estimation requires care in dealing with.

P20L14 I would find it more honest to say 'by 2 km and 600 m in the examples shown.'

P23 DOIs are an efficient means of locating papers and their inclusion in references makes life easier for readers.

P23L2 The page number of Boersma 2004 is D04311.

Strange though it may seem, I also have several comments on the review from J. Kgaran.

- I see no reason to show the reader the value of the cost function through each iteration. The precise path by which a retrieval achieves convergence is of minimal interest to anyone, unless the cost surface is very complex (and the author's previous papers provide no reason to suspect that). I also know of no operational retrieval scheme which retains such information. If you mean to compare the costs from the two schemes, that is complicated by the removal of the a priori from one retrieval, such that the costs are not directly comparable. (It occurs from their later comments that Kgaran may believe that the change in grid is done

after a single iteration of the retrieval. My understanding was that two full sets of iteration are completed. If the authors *are* breaking into the iteration scheme, I agree with Kgaran.)

- I agree that completely throwing out their a priori information is excessive.

- I don't see why it's relevant that two different observations have different profiles of information content with height. The example in Figs. 3-6 appears to have been chosen to demonstrate the utility of the proposed technique in difficult conditions, specifically a mid-atmospheric dry layer where uncertainties are large due to the small signal.

P20L12 I agree that this statement is overly general.

- I hadn't noticed that the Au threshold changed from 0.8 to 0.9 between papers. This change should be made explicit.

Fig.6 Error bars would indeed be helpful. Perhaps rendered as shading rather than whiskers (as in Fig. 10 of Sica and Haefele, 2015)?

P13L10 Yeah, it's the retrieval that has sensitivity (not the kernel) and, because you used the prior to make the retrieval grid, there will by a small, but non-zero, contribution to the eventual product so you should probably tone those statements down a bit.

- I don't agree that the large averaging kernel value is 'artificial'. It's not exact (in that a simultaneous observation would give subtly different values) but it's a fair guess of the information content. In essence, the technique is averaging observations until the SNR is large enough to permit a retrieval. These authors reduce the number of state vector elements to better match the quantity of information available. The idea is discussed (obtusely) in §5.8.3 of Rodgers (2000).

- I don't see why the trace of the averaging kernel would be 'artificial', since it's preserved from the fine resolution retrieval. Yes, the precise value is noisy but measurements are intrinsically imperfect; we can only do our best.

P12L15 The authors don't state that the method is poor in high SNR conditions. The second retrieval merely does no better than the first in such conditions. (Well, it's marginally worse since the resolution decreased.)

- I'm sympathetic with Kgaran that removing the a priori, by definition, cannot add information to the retrieval and therefore cannot increase the height of available information. I think that's being overly exact but isn't conceptually wrong.

P20L21 The uncertainty increases because they've decreased the information available. The resolution increases because it's a fairer representation of the information available. Their previous retrieval was unduly confident as the authors appear to believe there is no unbiased prior. Perhaps by adding a measure of the RMS vs. radiosondes, the authors can more clearly explain the merit of their adjusted data?

P9L22 I suspect this is the same point I made in the PDF, whereby in English 'increased resolution' means 'more fine detail' but the authors have used it to mean 'greater distance represented by each pixel.' That should be fixed throughout as it's quite confusing.

P19L12 Isn't the dry layer the area where the measured mixing ratio is substantially lower than the a priori profile?

P21L7 Actually, that's a good point. Removing the a priori is mostly going to convince inexperienced users that they don't need to think about it. The appropriate selection of prior information is the difference between a respectable retrieval and a black box that just provides the answer you wanted.

Please also note the supplement to this comment:
https://www.atmos-meas-tech-discuss.net/amt-2018-347/amt-2018-347-RC1-supplement.pdf

---