Response to RC1

We thank the reviewer for their feedback.

Response to RC2

We thank the reviewer for their constructive comments. We have responded to each comment in-line below, and will update the manuscript accordingly as described here.

Page 1: Line 2: by EarthCARE --> by the EarthCARE instruments.

Done

Line 4: and rain, and liquid clouds ; is the second and necessary there, I would remove it

Done.

Line 10/11 : you talk about three scenes but in line 11 say ‘using both case studies’. In the evaluation you only use the Hawaii, Halifax and Halifax-aerosol cases but not the third one. Please rephrase.

We agree that that was confusing. This now reads, “We have demonstrated and evaluated the ACM-CAP product as applied to EarthCARE test scenes simulated from numerical weather model forecasts, evaluating the product against the simulated measurements and the ``true'' quantities from the numerical model.”

Page 2: Line 51: clouds are now mature --> clouds can now be considered mature.

Thank you; done.

Line 55: , the Moderate Resolution Imaging… ; also add the before ‘Cloud and the Earth’s …..’

Thank you; done.

Page 3: Lines 58-59: Please reformulate the line ‘One advantage …. Scenes;

This now reads, “Single-instrument retrievals can be especially subject to uncertainties in complex or layered scenes---a limitation that multiple-instrument synergies can help to overcome.”

Line 59-60 Please reformulate line ‘For example ….phase clouds’ the second part does not seem to fit.

Thank you; done

Page 4: Line 105: within in the ESA (remove in)

Thank you; done

Line 117: Reference to Mason et al incorrect.

We've removed the reference; this paper hasn't yet been submitted

Line 122 for each. --> for each granule.

Thank you; done

Page 5: Line 124: Missing citation to AC-TC

Thank you; done

Page 6: Line 138: Missing citations to A-EBD, and in 140: to C-CD

Thank you; done

Lines 142: What seems to be missing from the paper is a bit of context as described in the paper of Eisinger. The JSG should be mentioned here by a bit more than just the name, but also the data throughput of the data processing from single instrument L2a --> synergistic L2b. Please add a paragraph for context on the EarthCARE processing chain and choice of JSG

We have added details about X-MET and X-JSG here, including links to the Eisinger paper. This paragraph now reads,

"CPR and MSI products are provided on their own instrument grids, while A-EBD is on the joint standard grid (JSG), which is initially defined by the auxiliary X-JSG data product \citep{Eisinger2022} and then inherited by all ATLID and downstream synergistic L2 data products. The JSG provides the common reference grid at 1km horizontal (along-track) and 100m vertical resolution onto which all active measurements are mapped. Within the ACM-CAP processor the CPR and MSI measurements are first interpolated onto the common grid before the retrieval is carried out. To inform the interpretation and assimilation of each measurement, additional variables describing measurement uncertainties, and quality and detection statuses, are also read from each data product."

Line 157: In practice both the Hessian…

Thank you; done

Page 7: Line 192.: On the discussion of why all variables are in elog space. The one issue which could arise is that when no negative retrievals (within the errors) are possible a bias could be introduced. Is this potentially the case, and if not why?

This is an interesting question. The first point is that an additional reason for using the logarithm is that many observed and retrieved variables, especially radar reflectivity, span many orders of magnitude; for this reason the use of logarithms is quite common and indeed the official CloudSat algorithms take the same approach (e.g. Austin et al., JGR 2009; Leinonen et al., JAMC 2016, which we now cite). But to fully address the question posed by the reviewer almost requires a separate study. Our understanding is as follows. By using logarithmic state variables, e.g. x=ln(LWC), the algorithm retrieves the expected value of x and an estimate of the standard deviation "sigma" of the probability distribution which is implicitly assumed to be Gaussian in ln(LWC) space. If we were to transform this probability distribution into linear LWC space then it would have a lognormal form such that the median of the new probability distribution is equal to exp(x), but the expected value would be larger with a value of exp(x+sigma^2/2) according to the property of lognormal distributions. If one wanted to know, say, the mean LWC of all clouds globally, then perhaps accounting for this difference would be needed, and it would be perfectly possible to compute given that we report the sigma value for all retrieved variables in the output product. But for radiative transfer applications, one of the major concerns of the EarthCARE mission, it is better to report the original retrieval, i.e. the expected value of the ln(LWC) probability distribution. This is because radiative properties of a cloud (shortwave albedo and longwave emissivity) are more linearly related to ln(optical depth) and hence ln(LWC) than they are to optical depth or LWC. These issues are applicable to a large number of radar algorithms, not just our own, so they deserve to be studied in detail in a future publication.

Table 1: This table is essential in order to understand how everything is linked and where the info comes from.

However, there is a problem with this as you describe parameters as, e.g. ln(S) but provide the linear value in the second column. Obviously I understand the why but it hurts the brain as it is confusing. My suggestion would be to show all parameters which are used logarithmically internally in boldface and remove the ln in the first columns. Describe this in the caption and all should be well.

The problem with Table 1, which we have now fixed, is that some of the stated a-priori values were for the logarithm of the microphysical quantity (e.g. ln(alpha_v) in ice), while others were of the quantity itself (e.g. LWC). We have now changed these entries so that all state the a-priori values for the actual state variable, i.e. the natural logarithm of the quantity (e.g. we now state "ln(10^-4)" for the a-priori of ln(LWC)). We do not think that using bold would be helpful since bold is used elsewhere in the paper to indicate vectors and matrices. So long as Table 1 is clear and consistent, and properly explained in the caption, we think this approach is preferrable.

The one remaining issue is the Lidar ratio, in this table you consistently use the backscatter to extinction ratio (units sr-1). However in the rest of the paper the lidar ratio is shown, in the more commonly used ratio, as extinction to backscatter ratio

(i.e. line 351). Please change the table definition of S, even if your code uses the inverse in reality.

One role of the paper is to act as documentation for the algorithm and the product, which works with backscatter-to-extinction ratio S, so it would create more confusion to make the paper inconsistent with the product. The reason that the retrieval itself uses backscatter-to-extinction ratio, rather than lidar ratio, is that the former is more natural particle scattering property being similar in concept to the single scattering albedo (i.e. the scatter-to-extinction ratio). Also, for a completely absorbing particle the former tends to zero while the latter tends to infinity. The reviewer is right that the paper is inconsistent though: we have now replaced some incorrect uses of the term "lidar ratio" with "backscatter-to-extinction ratio" (in the abstract and Table 1). In the results section 3.4 we now state that our retrievals will be shown in the form of lidar ratio (the reciprocal of backscatter-to-extinction ratio) in order to be consistent with other EarthCARE papers in this special issue. We now no longer inconsistently use "S" for lidar ratio in the results section with units sr^-1, but use the acronym LR with units of sr.

Page 11: Line 270: ‘is the retrieved as’ --> remove the

Thank you; done.

Line 293: Add references to A-TC and AC-TC

Thank you; done.

Page 12: Table 2 and Line 315 on page 13: Is Mie theory also used for dust aerosols? If so, what do you expect is the influence of this when used in the algorithm, do you experience biases with the MSI forward modelling due to this? Can you provide a short discussion on this around line 315

This table was inaccurate: we use lookup tables of the scattering properties of the four HETEAC species, wherein the non-sphericity of coarse dust is accounted for as described in the HETEAC paper (Wandinger et al 2022). HETEAC is now cited directly in the table, and the text has been updated to better describe the combination of lookup tables of scattering properties for the four HETEAC species.

Page 14: Line 333: ‘2 m surface winds’. You have not described where the 2m winds are used for. My guess is a Cox-Munk like approach for the SW calculations?

The surface winds were mentioned in error: we use the skin temperature for the thermal radiance model, but for the solar radiance forward model we rely on the surface albedo variable included in X-MET. Mention of the surface winds has been removed. An option for a Cox-Munk like approach making use of the surface winds is a good suggestion, and should be the subject of future work.

Page 17: Lines 422-425; Page 32 Lines 587-592. Both the forward modelling of the test scenes and those of ACM-CAP have started with the definition of the same four basic aerosol types defined for all wavelength by the HETEAC model and therefore span exactly the same parameter space. There is indeed no interaction with the cloud microphysics but neither is this the case in ACM-CAP, since the AC-TC defines the targets. Wouldn't it be more accurate to say that any differences seen between ACM-CAP and the model truth is could be due to differences in the RT models used and any LUTs resulting from this?

Thanks for pointing this out: these paragraphs certainly confused multiple issues, not all of which should not be interpreted as explaining the differences between the model truth and the retrieval.

The underlying HETEAC species in the test scenes and the retrieval are the same; however, the scene generation had access to the underlying quantities of the four HETEAC species freely varying in relation to one another, while the ACM-CAP retrieval can retrieve aerosols only with the predetermined mixtures of the four HETEAC species defined by the aerosol classes in A-TC. They may both be said to span the same parameter space, but the retrieval is not free to sample that parameter space evenly. The capacity of the ACM-CAP retrieval to recover the information in the model truth is therefore dependent on the classification from A-TC as well as the constraints applied within ACM-CAP. These latter indeed include the RT models and look-up tables, as well as the vertical representation and along-track Kalman smoothing, the treatment of observational errors, etc., etc.

We have re-written these two paragraphs to make fewer points more clearly.

Line 435: 'In the retrieved is systematically lower than the GEM model' --> The retrieved is systematically lower than the GEM model values

Thank you; done.

Line 439-442 'this Warm … temperature.' Rephrase this part, really hard to read, Warm --> warm

Thank you; done.

Line 445: sea level is a --> sea level shows a

Thank you; done.

Figure 3: Three areas of interest (A, B, C) are highlighted

Thank you; done.

Figure 12: panels a and b, text in bottom right not readable.

Thank you; done.

Page 29: tends --> tendency

Thank you; done.

Page 29, Line 545. Exhibit strong correlations.   I agree that these are very strong correlations, similar to 14b-c , 16c. Doesn't this just show how well the redistribution of information by the integrated measurements (MSI, PIA) is performed in the algorithm. The solution converges towards a fit of this information and since the integrated values have a relative smaller error compared to a single pixel lidar or radar value within the profile the correlation is very high. If the above argument is correct it would be good to reflect this in the text.

In Figure 11 we are evaluating a retrieval against independent data, so a high correlation indicates a skilful retrieval, but in Fig. 12 we are comparing the forward model to the observations, so a high correlation simply indicates that the retrieval behaved correctly by trying to best fit the observations. We have changed the text so as not to imply that the correlation in Fig. 12 indicates skill but rather just a retrieval that is behaving sensibly.

Page 30: Line 563 add a comment that in reality the ECMWF data will have different values to the truth as well (now with respect to the modelled fields).

Done; we added a note that this uncertainty due to errors in the atmospheric state needs further investigation.

Page 34: Line 621 which either is embedded

Thank you; done.