

Response to reviewer 1.

We thank the reviewer for taking the time and effort to comment on our paper and for their insightful and constructive comments. We respond below to each of their points in turn.

1. The abstract could also be more concise. Some of the background in the first paragraph could be left for the Introduction and some of the details in the 3rd paragraph could be consolidated into a more concise description of the results.

Response: We agree with the reviewer that some of the first paragraph could be moved to the introduction, and that the third paragraph could be more concise. The abstract has been revised along the lines suggested by the reviewer. Material from the original first paragraph of the abstract has been moved to the introduction, lines 14-36 in the revised paper, and the third paragraph has been condensed to focus on the main results.

2. The authors have decided to combine the description of CCREST-M and the algorithm evaluation into a single manuscript. While I think I understand the motivation for this choice, these topics could probably be covered in two distinct, concise, and equally interesting, papers.

Response: We thank the reviewer for their thoughtful suggestion and for indicating that they understand the motivation behind our choice. After careful consideration, we have decided to retain the integrated structure for the following reasons.

Firstly, the CCREST-M campaign was designed primarily as a testbed for ice-cloud scattering and radiative transfer studies. With the radar-radiometer geometry, flight patterns, and in-situ sampling strategy all configured to support the PSD retrieval and validation presented in a single paper. We feel this is clearer in one paper, rather than in several. The campaign description, its motivation, and the retrieval methodology are intrinsically linked. The rationale for the figure-of-eight overpasses with zenith-pointing radars, the post overpass in-situ descents, the use of dropsondes for atmospheric profiles, only becomes fully transparent in the context of the retrieval framework that they were designed to enable. Separating these into two papers would require duplication of context in say a data paper to motivate the measurement strategy.

Secondly, a standalone CCREST-M data paper say, would still require analysis of the data to demonstrate its utility and validity of the dataset through worked examples and analysis, which have already been presented in the present paper. Our understanding is that a data paper should not just be a catalogue, but a scientific contribution.

Thirdly, we believe that the integrated paper is more useful for the community, as future users of the CCREST-M data, whether for radar and/or radiative transfer studies would have the benefit of a single citable reference that documents the campaign, flight strategy, the instrumentation and data processing, and a validated dataset. Splitting these up into two papers would require users to cite and cross-reference both and risks the papers being cited without the methodological context in one of them.

We do, however, fully agree with the reviewer that the manuscript is dense and repetitive, we have addressed this concern through the changes detailed in our response to point 3 below.

3. The results (Section 5) are very thorough, but somewhat repetitive. For example, many figures are repeated for different variables, temperature ranges, and different cases and several points are reiterated multiple times. I recognize that each case illustrates a distinct aspect of algorithm performance, but I wonder if the results would be more impactful if they were conveyed more succinctly, perhaps consolidating the complementary findings from individual cases/comparisons where possible. Also, could some of the figures that represent variants of the same plot be moved to an appendix or supplementary material?

Response: We agree with the reviewer and have made substantial reductions to Section 5 along the lines suggested. The specific changes are as follows:

For the case C374, the original Figs. 9 and 10 have been consolidated into a single new Fig. 9, in which the ML-predicted moments assuming the gamma and exponential PSDs for the retrieved 3 GHz IWC are compared together with the in-situ moments. This consolidation has consequentially decreased the length of the corresponding sub-section accordingly. The original Figs. 11 and 14 have likewise been consolidated into a single new Fig. 11, which shows the comparison between the mean retrieved PSDs, assuming the gamma and exponential PSDs for the retrieval of IWC at 3 GHz, with the composite in-situ PSDs for C374. The original Tables 2 and 3, which compared the estimated IWCs with the in-situ IWCs have now been consolidated into a new Table 2. For the comparison between radar reflectivity simulations and the GRaCE 200 GHz observations we have retained Fig. 13, assuming both PSD assumptions, and removed the original Fig. 15 assuming the exponential PSD.

For C379, this section has been reduced by removing the original Fig. 18, which showed the measurement residuals $\ll 1$ dBZ, and so this is now just stated in the text, since the retrieval behaviour is no different to that already presented for the case C374. The figure presenting the moment comparisons in sub-section 5.2.1 on page 28 has been significantly improved in terms of resolution and text size. In replotting this figure, it was found that an outdated in-situ moment file had been used in the original submission. The corrected Fig 14 on page 28 of the revised manuscript shows that the PDFs of the ML-predicted moments overlap with the in-situ moments better than previously indicated. It has been verified that this correction is confined to the moment comparison figures and does not affect any of the other results presented in this sub-section.

For C382, in sub-section 5.3.1, the original Fig. 19 has been replaced by Fig. 16 in the revised manuscript on page 32, and as for C379, an outdated in-situ moment file had been used in the original submission. In this case, however, the corrected comparisons are not improved, owing to the failure of the ML method for this case as discussed in the manuscript. Again, this correction is confined only to the moment comparison figures. To reduce sub-section 5.3.2, the original Figs. 21 and

22 have been removed. The investigation that those figures supported, using the in-situ derived PSD parameters, is now described in the text alone, since this approach did not improve the results and the figures are not strictly necessary to convey the result. Figure 18 (originally Fig. 23) of the revised manuscript has been improved in terms of resolution and text size and retained, since it demonstrates that for this case the retrievals are stable, supporting our interpretation that the discrepancy between in-situ and retrieved PSDs is most likely owing to cloud evolution between the times of radar sampling and in-situ measurements.

Together, these changes have reduced the length of Section 5 significantly while retaining the essential scientific content.

4. The labels on several figures are very small (e.g. Figures 4, 5, 8, 9, and 10), some to the point of being illegible (11, 12, 14, 17, 20 - 23).

Response: We agree with the reviewer that the labels on the highlighted figures in the original submission were too small. All the flagged figures have been re-done in terms of label size (also ticks and legends), and resolution. Where figures have been consolidated or removed (as described in our response to point 3), this specific issue has been resolved by their replacement or removal.

Response to reviewer 2.

We thank the reviewer for taking the time and effort to comment on our paper and for their insightful and constructive comments. We are grateful that the reviewer finds the CCREST-M data valuable and has no critiques on the presented results. We respond below to each of their points in turn.

Major comments,

1. As discussed above, I am very confused that you did not use a triple-frequency approach. Alternatively, is it possible to use the Ka-band or W-band radar in the ML model? Then, you do not need the third frequency.

Response: The use of triple-frequency radar retrievals via say Bayesian optimal estimation is well established, and would not, in our view, justify a paper. We instead wished to explore an alternative, such as a data-driven approach, in which an ML ensemble model is trained to provide prior information on the cloud state, through both the thermodynamics and microphysics to inform a physics-based estimation of the PSD gamma parameters. The 3 GHz radar is required within this framework to provide the vertical profile of IWC, and is well suited for this purpose, since the 3 GHz radar return is in the Rayleigh scattering regime and is therefore directly related to the mass of ice in the column. The retrieval profile of IWC is then used as a feature to the ML ensemble, which ultimately provides first-guess profiles of gamma parameters, before the dual-frequency 35 and 94 GHz retrieval refines the PSD parameters. The third frequency is therefore not redundant in our framework but plays a distinct and necessary role. Our approach is largely supported by the findings, particularly for cases C374 and C379, and is further validated for C374 using the independent 200 GHz GRaCE radar.

The alternative suggestion made by the reviewer of using the Ka- and W-band reflectivity directly as a feature in the ML model is in-principle viable as yet another alternative approach. However, we do not have access to a Chilbolton dataset that was specifically linked to the PICASSO campaign. Also, the PICASSO campaign did not profile the whole cloud, which is needed in our case for future radiative transfer studies. NASA's IMPACTS campaign is where such a dataset might be found, but this is an entirely different location to the Chilbolton Observatory. Such an alternative is beyond the scope of our study and paper. We leave this for others to pursue. Furthermore, we deliberately wished to use ML as a first guess to inform the physical retrieval, rather than the retrieval method itself, the physical retrieval anchored on the scattering model provides a layer of physical constraint that we consider important.

2. The adequacy of using a single ice type. It seems to be a bold assumption to me that a single ice type was used in the radar retrieval. I understand that the CPI imagers suggest the presence of ice rosettes, but the sampling area is very limited compared to radar observations. I would encourage a thorough discussion on this limitation.

Response: We thank the reviewer for raising this point. We would first emphasise that the CIP-15 imagery, while limited in sampling volume relative to the radar, is consistent across all three of the in-situ sampled cases (Fig. 3) and is consistent with the established literature on mid-latitude ice cloud microphysics (Lawson et al., 2019, Wagner et al. 2025, and references therein). The consensus is that within in-situ generated ice cloud, rosettes and forms of rosettes comprise more than 50% of the ice crystals, see for instance Wagner

et al. (2025) on the prevalence of bullet rosettes within in-situ generated ice cloud. Moreover, since we restrict our analysis to particle sizes greater than 100 μm , the occurrence of budding or growing/grown aggregated rosettes are even more common as shown in Lawson et al. (2019).

We would also like to clarify a point that we recognize was not made sufficiently explicit in the original submission. The rosette-aggregate scattering model should not be thought of as a single ice type, but rather a morphologically diverse ensemble. As described in the original manuscript and in detail by Kleanthous et al. (2024), the aggregates are constructed from solid three-branched rosette monomers and largely grow into different aggregate forms as size increases, with the mass being constrained by the observed Cotton et al. (2013) mass–dimension relationship.

Finally, as pointed out in the original manuscript, the rosette aggregate model has previously been tested against triple-frequency radar observations from NASA’s IMPACTS campaign, which also took place over mid-latitude regions, but off the coast of the eastern USA. In Baran et al., 2024, as cited in the paper on page 11, lines 292-296, of the original manuscript, showed that the model reproduced the observed reflectivities at 9, 35 and 94 GHz to within a few dBZ across four IMPACTS case studies. The applicability of the model to mid-latitude frontal ice cloud is therefore supported both by the in-situ habit identification for CCREST-M, the 200 GHz simulations presented in the paper, and by independent triple-frequency validation in a separate mid-latitude campaign in a different mid-latitude region.

To clarify the morphologically diverse nature of the rosette-aggregate model, we have added the following sentence in the revised paper on page 11, lines 288 – 291, we state:

“Although referred to as the rosette-aggregate model, this representation does not correspond to a single ice type. As detailed in Kleanthous et al. (2024), the aggregates are composed of rosettes but evolve into different structural forms as they grow, with morphology varying systematically when mass scales with the square of their maximum dimension.”

The Wagner et al. (2025) citation has been added to the reference list in the revision, and cited on page 11, line 286.

3. Redundant figures for PSD comparisons. It is recommended that in-situ observations, ML predictions and dual-frequency radar retrievals be integrated into a single panel. In this way, direct and intuitive comparisons among PSD results from different methods can be easily conducted.

4. Similarly, Tables 2,3,4 should be integrated into one table for a direct comparison. The same to all the PDF plots of different moments (e.g., fig9&10).

Response: We thank the reviewer for these suggestions, and as Reviewer 1 raised a similar point about the size of the manuscript, we treat comments 3 and 4 together here. We have made substantial changes to Section 5 along the lines suggested. The specific changes are as follows:

For the case C374, the original Figs. 9 and 10 have been consolidated into a single new Fig. 9, in which the ML-predicted moments assuming the gamma and exponential PSDs for the retrieved 3 GHz IWC are compared together with the in-situ moments. This consolidation

has consequentially decreased the length of the corresponding sub-section accordingly. The original Figs. 11 and 14 have likewise been consolidated into a single new Fig. 11, which shows the comparison between the mean retrieved PSDs, assuming the gamma and exponential PSDs for the retrieval of IWC at 3 GHz, with the composite in-situ PSDs for C374. The original Tables 2 and 3, which compared the estimated IWCs with the in-situ IWCs have now been consolidated into a new Table 2. For the comparison between radar reflectivity simulations and the GRaCE 200 GHz observations we have retained Fig. 13, assuming the gamma and exponential PSDs, and removed the original Fig. 15 assuming the exponential PSD.

For C379, this section has been reduced by removing the original Fig. 18, which showed the measurement residuals $\ll 1$ dBZ, and so this is now just stated in the text, since the retrieval behaviour is no different to that already presented for the case C374. The figure presenting the moment comparisons in sub-section 5.2.1 on page 28 has been significantly improved in terms of resolution and text size. In re-plotting this figure, it was found that an outdated in-situ moment file had been used in the original submission. The corrected Fig 14 on page 28 of the revised manuscript shows that the PDFs of the ML-predicted moments overlap with the in-situ moments better than previously indicated. It has been verified that this correction is confined to the moment comparison figures and does not affect any of the other results presented in this sub-section.

For C382, in sub-section 5.3.1, the original Fig. 19 has been replaced by Fig. 16 in the revised manuscript on page 32, and as for C379, an outdated in-situ moment file had been used in the original submission. In this case, however, the corrected comparisons are not improved, owing to the failure of the ML method for this case as discussed in the manuscript. Again, this correction is confined only to the moment comparison figures. To reduce sub-section 5.3.2, the original Figs. 21 and 22 have been removed. The investigation that those figures supported, using the in-situ derived PSD parameters, is now described in the text alone, since this approach did not improve the results and the figures are not strictly necessary to convey the result. Figure 18 (originally Fig. 23) of the revised manuscript has been improved in terms of resolution and text size and retained, since it demonstrates that for this case the retrievals are stable, supporting our interpretation that the discrepancy between in-situ and retrieved PSDs is most likely owing to cloud evolution between the times of radar sampling and in-situ measurements.

Together, these changes have reduced the length of Section 5 significantly while retaining the essential scientific content.

5. ML predictions for PSD moments were validated to in situ observations, but the dual-frequency validations are missing.

Response: The dual-frequency radars are used to retrieve the gamma PSD parameters, and it is the shape of the resulting PSDs that is compared with the in-situ PSD shapes. Comparisons in moment space would yield equivalent information, since the moments are derivable from the PSD shapes. We prefer the PSD-shape comparison because it is the PSDs themselves that will be used as inputs to the radiative transfer simulations in a subsequent paper, and it is the convolution of the PSD-shape with the scattering properties that determines the simulated brightness temperatures. This rationale is briefly stated in the original manuscript on page 23, line 585. In addition, the independent validation against the 200 GHz GRaCE radar, which the reviewer commented positively on

below, is particularly relevant, as 200 GHz lies in the mm-wave spectrum of interest to the subsequent radiative transfer paper.

6. I like the G-band validation part. Since the G-band radar was collocated with other radars, I would recommend a long-term validation. I believe it is very handy to implement.

Response: We thank the reviewer for the positive comment on this aspect of the paper, and for the suggestion. We agree that the G-band radar provides good independent validation of the retrieved PSDs and scattering model. We were fortunate that the G-band radar was operating during the C374 case. The G-band radar at Chilbolton is, however, an experimental radar and is not run routinely. It acts as a proof-of-concept for a future G-band radar in space. It has been operated during several other campaigns besides CCREST-M. Given that the radar is experimental, the suggestion made by the reviewer is not feasible, and so is beyond the scope of the present paper.

Some technical comments,

1. Figure 2. Mark the periods where the aircraft validations were made.

Response: We thank the reviewer for their suggestion. The aircraft periods of radial legs, figure-of-eight manoeuvres, and in-situ sampling have now been marked on a new Fig. 2 in the revised paper.

2. L204. Liquid clouds are not uniformly distributed. How did you do the liquid attenuation correction with LWP?

Response: We thank the reviewer for this point. We assume that the liquid water cloud beneath the ice cloud is effectively a single layer, and we assign a representative in-cloud temperature for calculating the specific liquid water absorption, which we then multiply by the HATPRO LWP to obtain the one-way attenuation. This has been clarified in the revised paper on page 7, lines 204-209, in the revised paper with the following text:

“ To account for liquid water cloud attenuation, cloud liquid water path (LWP) was retrieved from the Chilbolton HATPRO microwave radiometer (Walden, 2026a), with missing data interpolated over short time gaps. The HATPRO LWP provides a column integrated value and does not constrain the vertical distribution of liquid water. We therefore assume the liquid cloud to be a single layer, with a representative in-cloud temperature taken from the interpolated atmospheric profile at the liquid-cloud height, where the cloud top height was manually defined using the ceilometer (Walden, 2026b), lidar, and radar backscatter profiles. The specific liquid water absorption at each radar frequency was computed from the Ellison (2007) absorption model at this representative temperature and multiplied by the HATPRO LWP to give the one-way attenuation.”

3. L212. L267. Looks conflicting. You definitely need gaseous attenuation for Ka- and W-band radars.

Response: We thank the reviewer for noting this contradiction. Agreed, the 35 and 94 GHz observations have now been corrected for gaseous attenuation. We have therefore replaced the original discussion with the following text in the revised paper on page 8, lines 214 – 224, this reads as follows:

“To account for gaseous attenuation by oxygen and water vapour on the radar reflectivities at 35 and 94 GHz, we computed the two-way partial-column gaseous attenuation from near-surface to each in-cloud retrieval gate using the aircraft dropsonde-derived profiles of temperature, pressure and water vapour volume mixing ratio released near the Chilbolton Observatory. A simplified gas absorption parametrisation was used, consisting of pressure-scaled oxygen absorption and humidity-scaled water vapour absorption, with coefficients providing representative specific attenuation values of 0.02, and 0.1 dB/km for oxygen, and 0.01, and 0.08 dB/km for water vapour at 35 and 94 GHz, respectively, under standard atmospheric conditions as recommended by ITU-R P676.13 (2022). The mean partial-column two-way attenuation at 94 GHz reaches approximately 1.5 dB at 94 GHz for C374, this is consistent with Hogan et al. (2000), who using co-located 35 and 94 GHz radars at Chilbolton Observatory reported a two-way gaseous attenuation at 94 GHz of approximately 2 dB from near-surface to 10 km. At 3 GHz, gaseous attenuation is negligible (specific attenuation values of order 0.001 dB/km for oxygen and zero for water vapour) and is not considered further. ”

4. L210. It should be 1 dB. In addition, I would not say ‘much less’ than 1 dB. We recently compared different parameterizations, and it is sub-dB difference. Li, Q., Li, H., Sun, X., et al. (2026). A survey of snow growth signatures from tropics to Antarctica using triple-frequency radar observations. *Atmospheric Chemistry and Physics*, 26(2), 1249-1264.

Response: We thank the reviewer for the pointer to their recent multi-frequency and multi-campaign paper. The actual values quoted by McCusker et al. (2024) is $\ll 1$ dBZ. In the revised paper, On page 7, lines 210-214, we re-word the original submission to be:

“It was previously found by McCusker et al. (2024) that for similar frontal iced-cloud conditions the two-way path integrated ice attenuation at these frequencies to be well below 1 dBZ. This is consistent with the recent multi-campaign comparison by Li et al. (2026), who reported a median W-band path-integrated snow attenuation of 0.3 dB. Therefore, ice attenuation is not corrected for explicitly in this paper.”

The Li et al., 2026 citation has been added to the reference list.