Review of manuscript #AMT-2021-111 entitled 'Spaceborne differential absorption radar water vapor retrieval capabilities in tropical and subtropical boundary layer cloud regimes' by Richard Roy, Matthew Lebsock, and Marcin Kurowski.

The manuscript provides a detailed evaluation of the potential of a differential absorption radar (DAR) operating in three tones around the 183 GHz water vapor absorption line for PBL humidity profiling inside clouds, precipitation and for measuring integrated water vapor (IWV) in clear air regions. While the concepts presented by the authors are innovative, I believe the study overlooks a few effects that would have a considerable impact on the results (see major comments below). Given this, my recommendation is to reject the manuscript in its current form and invite the authors to submit a revised manuscript when these effects are simulated.

Major comments:

1. I believe the authors have omitted to consider the spaceborne range weighting function in their forward simulations. Past literature has shown that the range weighting function of long pulsed radars effectively stretches along, or in other words vertically smooths, any vertically-narrow cloud or precipitation feature. I believe this would affect the forward simulated radar observables that the authors present and rely on to assess their technique. Beyond vertically stretching/smoothing the fields, the range weighting function should be expected to introduce covariance of the random error. The authors correctly introduce random error after the retrievals are performed based on the linearity of the transfer function from the state vector to observations (Page 15); however, I believe this only represents one source of the true error, which should also include error caused by the covariances along range especially if the observations are to be oversampled by a factor of 4 of the true radar resolution to achieve the desired range gate spacing.

   The reviewer is correct that effects on the water vapor retrievals coming from a specific range weighting function (RWF) are not included in this work. While we will address the individual points mentioned above, we first note that our approach is equivalent to assuming a RWF that is a top-hat function of width 50 m. This is useful from a measurement capability assessment point of view because it does not assume specific implementation-level aspects of the transmit waveform and pulse compression filter. Another way of specifying the RWF properties that would be applicable to this study is to assert that, regardless of the chosen waveform parameters, the RWF must ensure negligible random error covariance between range bins that are separated by 50 m. While this is admittedly idealized in some respects, we still find it useful for the purposes of this study as we're not interested in analyzing a particular implementation-specific architecture, but rather examining the fundamental limits of the DAR technique for PBL humidity observations.

   Another point to clarify before addressing the main point of the reviewer's main comment here is the effect of range oversampling of the retrieval. The resolution that is oversampled here is the water vapor profile resolution of 200 m, which is not inherently related to the range resolution or RWF. In terms of the radar signal processing, we are assuming that the radar echo power vs. range profile is sampled at exactly the range resolution of 50 m (i.e., neither undersampled nor oversampled). Therefore, it should NOT be thought of as oversampling the RWF by a factor of 4, which would lead to significant bin-to-bin covariance terms, but rather adequately sampling the RWF while purposefully downsampling when performing the water vapor retrieval.

To begin addressing the reviewer's main point, we agree that the effect of an RWF on long-duration-pulse radar observations is to stretch/smooth features in the vertical, where the relevant length scale for the "smoothing" is the width of the RWF itself. The degree to which this smoothing will become problematic for the retrieval requires comparing the characteristic (vertical) length scale over which the cloud/precipitation reflectivity profiles change appreciably with the width of the RWF. This is of course a detailed task that depends sensitively on both the cloud/precipitation regime and the implementation-specific RWF, and so we don't believe that it belongs in this manuscript which is focused on addressing higher-level questions of the spaceborne DAR capabilities. However, we do provide in this response a detailed assessment of the RWF impacts on retrieval biases for a specific profile from the RICO case to highlight that, with sufficient pulse chirp bandwidth, the impact of bin-to-bin correlations on the retrievals is minimal. This discussion is provided here for the reviewers, however we do not believe that this material is appropriate for the manuscript.

We begin our RICO RWF case study by specifying a few candidate transmit waveforms, and by using the matched filter as the pulse compression filter. To match the notional high-level parameters used in the main manuscript, we assume a pulse length of 50 us, a Hann window for the transmit waveform taper, and a list of candidate chirp bandwidths: B = 0.75, 1.5, 3, 6, and 12 MHz. We then calculate the RWF, as well as two parameters related to the width of the RWF: (1) the Fourier limited resolution in range, c/(2B); and (2) the effective range resolution calculated as the equivalent width of a top-hat RWF with the same area as the true RWF. This latter range resolution (2) is more representative of the true range resolution of the system, and is exactly the range resolution term that arises in the weather radar (or volumetric scattering radar) equation. The plot below describes the difference between these two parameters, plotted against the compressed pulse time axis, which can be converted to range by multiplying by the factor c/2.

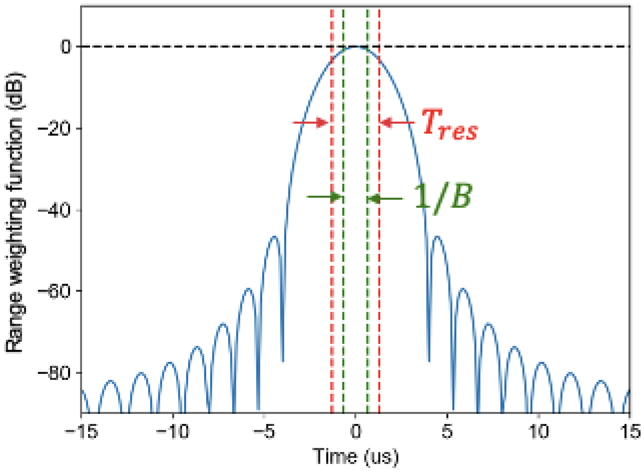

Below we plot the individual RWFs for each candidate chirp bandwidth, as well as provide a table of the relevant range resolution parameters.

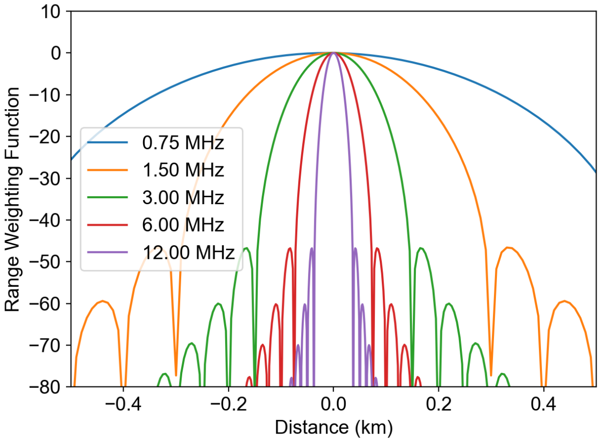

B (MHz)	c/2B	$cT_{res}/2$
0.75	200 m	389 m
1.5	100 m	194 m
3	50 m	97 m
6	25 m	49 m
12	12.5 m	24 m

The next task is to create simulated reflectivity profiles for each RWF (and each of the 3 DAR channels) and perform a retrieval in each case. Here we do so by working in the single-scattering approximation, and by convolving the simulated, attenuated reflectivity profile (which is at the very high LES vertical resolution) with the RWF, and downsampling the resulting convolved profile at the $cT_{res}/2$ effective range resolution (Shown in Figure below). Then, we perform the water vapor retrieval using the treatment outlined in the “DAR Measurement Methodology” Section of the manuscript (new Section 2.3). Examples of the convolved reflectivity profiles before downsampling are shown below for the 155.5 GHz DAR channel:

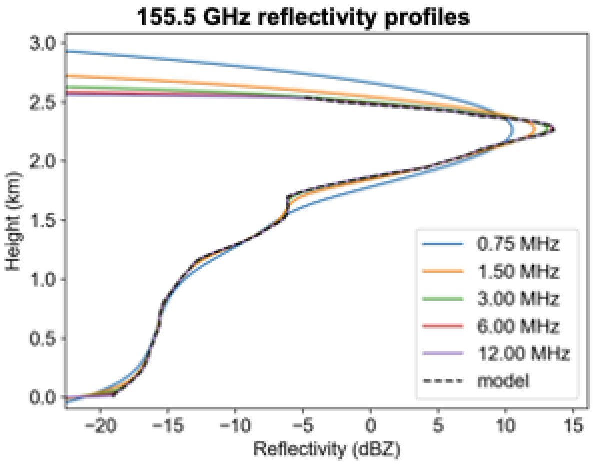

Note that we have chosen a profile from RICO that contains a sharp cloud/precipitation feature just above 2 km, which is the type of feature that could result in RWF-induced biases in the water vapor retrieval for coarse enough range resolution (i.e., small chirp bandwidth). Furthermore, example downsampled (solid circles) profiles at the 3 DAR frequencies for 2 different chirp bandwidths are shown below:

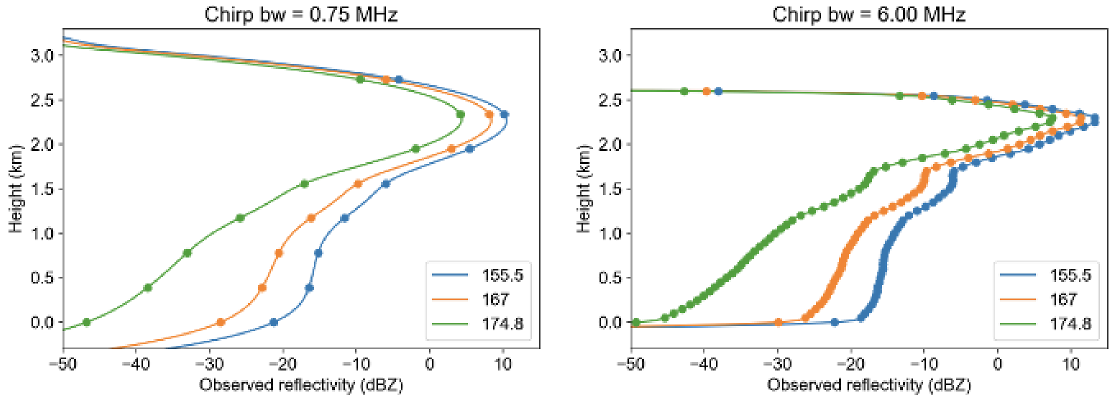

Note that we do not add random error/noise to these profiles, though as the reviewer highlighted it can in theory be introduced after the retrievals are performed based on the linearity of the transfer function from the state vector to observations. Instead, what is most important to assess from a bin-to-bin covariance point of view is the impact of the RWF convolution on retrieval biases (i.e., systematic error). The resulting retrievals for each chirp bandwidth are shown below, where we've used a retrieval step size of r^{+}-r^{-} = 200 m (see Section 2.3):

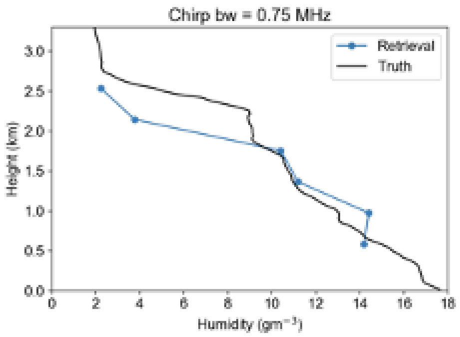

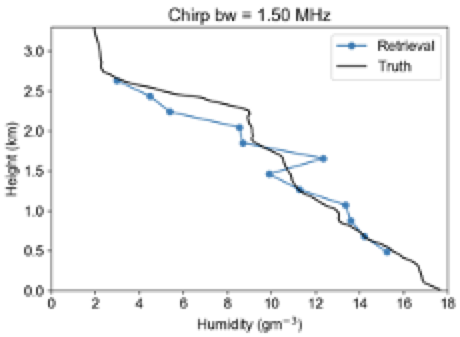

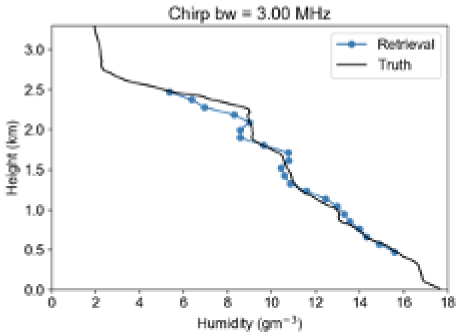

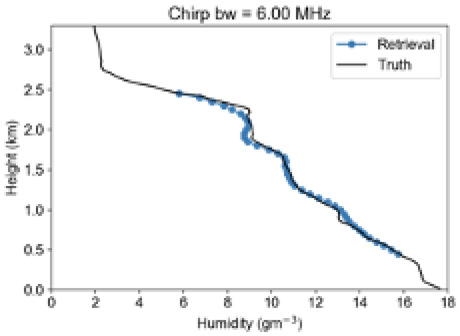

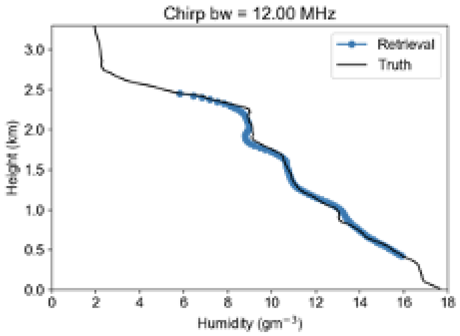

As expected, the low chirp bandwidths of 0.75 and 1.5 MHz, which either undersample (0.75) or provide commensurate sampling (1.5) of the retrieval resolution of 200 m, result in significant biases, especially in the vicinity of the strong reflectivity feature above 2 km. However, for larger chirp bandwidths, the stretching effects of the RWF convolution do not affect the retrieval accuracy. Importantly, the range resolution assumed in our simulation study in the manuscript of 50 m corresponds to a chirp bandwidth of 6 MHz, which is well into the regime where the RWF stretching does not affect the retrieval.

This case study shows that, even for a pessimistic scenario where there are sharp features in the reflectivity profile, the assumed radar range resolution of 50 m is consistent with a 6 MHz chirp and resulting RWF that limits bin-to-bin covariance to a level where it does not significantly impact the retrieved humidity accuracy. While this does not prove the independence of our analysis and results on the specifics of the RWF for all possible measurement scenarios and cloud/precip regimes, as there could be contrived examples where even the 6 MHz RWF results in bias issues, we believe it adequately addresses the reviewer's concern around the stretching from the RWF.

We do note here that, in the original manuscript on pg. 3 line 22, we state that a range resolution of 50 m is consistent with a pulse bandwidth of 3 MHz, which was incorrect. This is only the case for the c/2B figure, and for the purposes of the radar volumetric scattering sensitivity, it is more accurate to use the $cT_{res}/2$ figure. Therefore, in the revised manuscript we have edited this to read 6 MHz. Additionally, we have added a sentence immediately after this line acknowledging that these results do not include specific aspects of the RWF, and that our approach is equivalent to assuming a top-hat RWF with a width 50 m. Finally, we have added a sentence in the conclusions section (end of the first paragraph) noting that future investigation is needed to assess the impact of these biases for specific instrument implementations.

Added Text Page 3 Line 28: "We note that the approach implemented here does not include specific aspects of the radar range weighting function (RWF), and is instead equivalent to assuming an RWF that is a top-hat function of width 50 m."

Added Text Page 28 Line 1: "Future investigations focused on specific instrument architectures and radar waveform parameters are necessary to fully assess the impact of range-bin-to-range-bin correlations introduced by the range weighting function."

2. The authors discuss the use of a spaceborne DAR system operating a long pulse with frequency modulation like that used in RainCube. Based on my limited knowledge of RainCube's performance, I would expect that such a radar would suffer from contamination by the surface echo and range side lobes in the lowest kilometer of observations above the surface. Yet the authors failed to mention anything about how the proposed radar would be impacted by these two effects (surface echo and range side lobes) and about how this may impact the sensitivity of this radar to low level targets and the overall performance of the proposed retrieval in the boundary layer.

   The reviewer is correct that radars using pulse compression with large time-bandwidth products will generally have to contend with all of the impacts associated with the range weighting function. In the preceding major comment, we addressed this in the context of range-bin-to-range-bin covariance and potential impacts on the water vapor retrieval. As the reviewer points out, another potential issue stemming from the RWF is the surface response and its impact on the radar's sensitivity, especially in the lowest layers of the atmosphere. The surface response involves the combination of the beam pattern (with its own sidelobes), the RWF, and the projection of the beam onto a spherical surface, which results in surface returns that are spread over range. For a system like RainCube, this generally limits the detection of precipitation to heights above 500 m (see discussion in Battaglia et al., Spaceborne Cloud and Precipitation Radars: Status, Challenges, and Ways Forward, *Reviews of Geophysics*, 2020), though this neglects the added sidelobe suppression from beam attenuation in liquid. However, the surface clutter for the notional radar described here is greatly reduced relative to RainCube for two main reasons: (1) the chirp bandwidth in our case is (6 MHz)/(2.5 MHz) = 2.4 times larger; and (2) because of the shorter pulse duration relative to the pulse coherence time. In our case, we use a pulse duration that is around half the "time to independence", whereas in RainCube the pulse time of 166 us is significantly longer than the time to independence of around 25 us.

The preceding discussion was intended to provide context for why the situation here is much different than the only reference case for a long duration pulse radar (Raincube). To quantitatively assess the difference, it is necessary to implement a surface scattering model that includes a realistic beam pattern, surface illumination geometry, Doppler range migration, and the RWF. While we do not believe this level of detail belongs in the manuscript, we provide such a calculation here in response to this particular point. The surface clutter profile below assumes a realistic beam pattern resulting from the tapered illumination of a parabolic reflector with a Gaussian feed horn. For these simulations, we use the taper value of 10.9 dB that maximizes aperture efficiency. Furthermore, we use the same pulse parameters described above: 50 us pulse with 6 MHz chirp bandwidth and a Hann window. Guided by recent aircraft studies with JPLs Vapor In-cloud Profiling Radar (VIPR) instrument, we assume a surface normalized radar cross section (NRCS) of 15 dB, which is a pessimistically large value corresponding to very low wind speed conditions at near-nadir incidence. We perform the calculation at the middle DAR channel frequency of 168 GHz. Furthermore, we note that if we use the high-level parameters from the RainCube instrument in our surface clutter simulation (0.5 m aperture, 166 us pulse, 2.5 MHz chirp bandwidth), we reproduce the surface clutter suppression (~65 dB at 0.5 km above the surface) as modeled in Beauchamp et al., Observations and Design Considerations for Spaceborne Pulse Compression Weather Radar, *IEEE TGRS*, 2020 (https://ieeexplore.ieee.org/document/9169707).

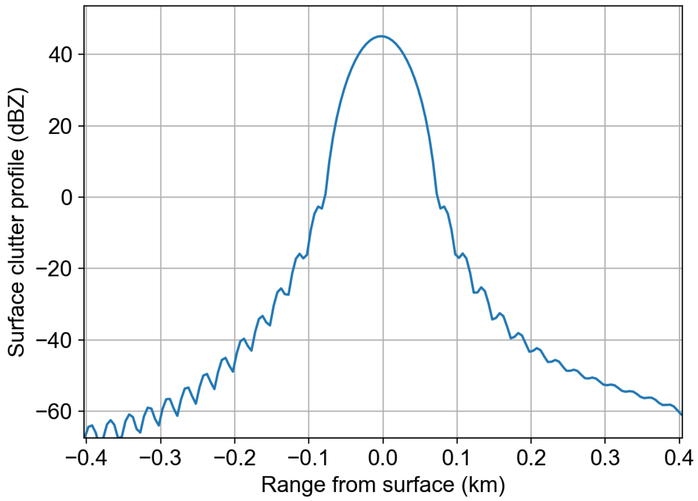

In the plot above, negative range corresponds to the atmospheric region of interest. We find that the clutter profile intersects the minimum detectable reflectivity (see Table 2 of the revised manuscript) of -34 dBZ just below 200 m above the surface. Therefore, we once again see that the large chirp bandwidth of 6 MHz (and corresponding fine range resolution of 50 m) is necessary to mitigate a variety of factors, including surface clutter, that could potentially limit DAR water vapor retrieval sensitivity and coverage. In short, surface clutter is not a limiting factor for the studies presented here because of the very high range resolution used.

We have added a short discussion of surface clutter at the end of the "Radar Forward Model" Section.

Added Text Page 8 Line 5: "In addition to thermal noise, surface clutter can potentially limit the radar detection sensitivity in the lowest layers of the atmosphere. In this case, we find that the large pulse bandwidth of 6 MHz and pulse duration that is well below the decoherence time result in a surface clutter profile that is largely confined to the lowest few hundred meters above the surface, with the surface equivalent reflectivity falling below -40 dBZ at a height of 200 m. Therefore, we do not expect this clutter source to limit detection sensitivity. We note that this contrasts with the experience of RainCube, where measurement sensitivity is degraded in the lowest 0.5-1 km of the atmosphere because of surface clutter (Battaglia et al., 2020). However, key differences between the two systems that explain this difference are the relative chirp bandwidths and the ratio of pulse duration to coherence time, which in the case of RainCube is larger than 1."

3. While it is interesting to consider the use of intelligent sampling using a passive/active sensor synergy, I believe the authors need to discuss more of the key aspects of such an endeavor to prove that it would even be feasible. For one, the authors did not discuss how a 2-m dish antenna will be configured to perform intelligent pointing. Overall, I believe a lot more information needs to be presented regarding this theoretical intelligent scanning technique. I would even argue that it should be the topic of a separate study.

   Intelligent scanning of cloud radars towards high-value targets is not a new idea, and there are several methods that could be technologically feasible for accomplishing this. The most obvious antenna solution that has airborne heritage (e.g., AirMASTR) is a phased-array fed reflector (PAFR), which in this case could consist of a line feed illuminating a 1D parabolic reflector. Such a system would easily be able to point in an agile fashion across the field of regard discussed in this work. Another potential solution is a mechanically steerable sub-reflector that illuminates a primary slightly off the paraboloid axis. This technique is routinely used in quasioptical systems with parabolic reflectors to achieve steerability. However, we believe that these details are beyond the scope of the current work, which is primarily concerned with answering the following question: If such idealized radar performance can be assumed, what is the quality of the water vapor retrievals that result from such a system in these various scenarios? We agree with the reviewer that the details of this implementation ought to be the topic of a separate study, and there are many implementation-specific problems to be solved. However, this paper is specifically not focused on the implementation-specific problems, and more on the basic physical limitations of DAR water vapor retrievals from space.

Minor comments:

1. P2, L14: The authors claim that they performed detailed orbital simulations. Please elaborate on what detailed information about the satellite orbit is used in this study beside the satellite altitude and velocity.

   For clarity we have changed this to ‘detailed radar simulations from an orbital altitude.’

2. P3, L31: I think the word “Table” is missing in front of “1”.

   The word “Table” was indeed missing. We have inserted this in the revised manuscript.

3. P6, L1: A formal study has been published based on Tatarevic et al., 2019 report. As such, I believe it is more appropriate to cite Oue, M., Tatarevic, A., Kollias, P., Wang, D., Yu, K., and Vogelmann, A. M.: The Cloud-resolving model Radar SIMulator (CR-SIM) Version 3.3: description and applications of a virtual observatory, Geosci. Model Dev., 13, 1975–1998, https://doi.org/10.5194/gmd-13-1975-2020, 2020

   Thank you for pointing us to this reference. We have replaced the reference to the report with the formal study reference.

4. P7, Fig. 1: Why did the authors not consider the Self-Similar Rayleigh-Gans Approximation (SSRGA) for the snow radar forward scattering modeling. Here is a useful reference: Hogan and Westbrook, 2014: Equation for the Microwave Backscatter Cross Section of Aggregate Snowflakes Using the Self-Similar Rayleigh-Gans Approximation, Journal of the Atmospheric Sciences, 71(9), 3292-3301.

   We are aware of the use of the self-similar Rayleigh-Gans approximation for modeling mm-wave scattering from snow aggregates, and there is no fundamental reason why this could not have been used in our study. However, SSRGA only applies to large aggregates consisting of many individual monomers since it relies on a statistical/ensemble-based decomposition of the particle mass distribution in Fourier space. Therefore, it would not apply to the pristine “ice” species in this work. We note that the main purpose for treating ice and snow scattering in this work is to assess the degree to which these particles introduce a bias in the DAR retrieval due to frequency-dependent scattering (i.e., that which is associated with the parameter a_2 in Eq. (10)). Specifically, note that we are not attempting to model the G-band radar’s ability to retrieve ice microphysical parameters, we simply want an approach that provides a realistic quantification of the frequency-dependence of the scattering. Since SSRGA is derived assuming an ensemble mean approach, there is good reason to believe that non-trivial frequency dependence that might be present in the single crystal realization case may be “washed out” by performing the ensemble average, but we have not done a detailed study of this (well beyond the current scope). It could be interesting in future work to investigate the different DAR biases predicted using the SSRGA and DDA formalisms, but this is certainly beyond the scope of the current work, where we’re primarily focused on PBL water vapor and therefore liquid hydrometeors.

5. P9, Fig. 2: Is it possible that the precipitation and cloud mixing ratio labels are mixed up in the right panel of figure a? I find it difficult to believe that the rain is elevated and higher in magnitude than the cloud mixing ratio.

   The labels in the figure are correct. A potential explanation is the presence of shear near cloud edges and the fact that rain has its own terminal velocity, which can separate it from the cloud regions. For such conditions, turbulent eddies can transport rain water away from the cloud at larger distances because its lifetime is longer than for cloud water in unsaturated regions. Cloud water evaporates quickly in unsaturated environments due to the bulk saturation approach, which is commonly used in LES modeling as an approximation of phase changes for smaller droplets. Rain water can exist longer because the evaporation rate is slower.

6. P9, Fig. 2: I would recommend not showing the water vapor retrievals before presenting section 2 that describes the retrieval.

   The water vapor retrievals shown in this figure do not use the new retrieval algorithm detailed in Section 2.4, but rather the simplified humidity retrieval algorithms discussed in 2.3 (Eqs. 9 and 10). That is, this figure is positioned within the same section that describes the algorithms used, and therefore we deem it appropriate. We have added a reference to the appropriate section 2.3 in the figure caption to clarify this.

7. P16, Fig. 3: The caption for panel d is missing.

   Thank you for catching this missing caption. We have corrected this in the revised manuscript.

8. P18, Fig. 5: I believe the surface echo and range side lobe should be visible in this figure. Afterall this is a chirp radar and its sensitivity is very high (better than CloudSat) and as such I would expect contamination of at least all the sub-cloud layer.

   Please see response to Major Comment 2 above. In short, we treat the surface response with associated range sidelobes as a height-dependent noise floor, which is subtracted off of the detected power vs. range profiles to arrive at a simulated observation of reflectivity (which should have no contribution from the thermal noise floor or from range sidelobes). Because of the high chirp bandwidth of 6 MHz, these range sidelobes are limited to the first 200 meters above the surface. This is one of the main reasons for choosing such a high range resolution - i.e., to limit the impact of range sidelobes on PBL observations.

9. Can you comment on the ability of the technique to capture the mesoscale organization of the humidity field?

   The main limitation in a DAR's ability to capture this sort of organization in the humidity field is the requisite horizontal averaging necessary to reduce measurement uncertainty to a sub-variability level. Practically, we don't see this as being possible for the profiling measurement. If one is focused on the total column water vapor retrieval, which can achieve low uncertainty with much less horizontal averaging, such a study could be feasible. However, it likely requires a full cross-track scanning system to resolve the 2D humidity field, which is not the concept considered here.

10. P21, F9: I believe the stretching effect of the range weighting function should be visible in this figure. A useful reference to visualize this effect is Lamer, K., Kollias, P., Battaglia, A., & Preval, S. (2020). Mind the gap–Part 1: Accurately locating warm marine boundary layer clouds and precipitation using spaceborne radars. AMT, 3(5), 2363-2379.

    Please see response to Major Comment 1. In short, our approach is equivalent to assuming a range weighting function (RWF) that is a top-hat function with a width given by the range resolution of 50 m. All spaceborne radars deployed to-date have utilized pulse parameters that result in much broader range weighting functions (and therefore range resolution), and so for those systems you would expect to see significant "stretching" or "blurring" of localized cloud features like these.

    In terms of justifying the top-hat RWF assumption, for the scope of questions we are interested in addressing in this work, we believe this is the best approach since it does not require a specific choice of radar waveform parameters (e.g., waveform taper/window), though we do provide some notional ideas of things like transmit power and chirp bandwidth. In reality, a mission would go through a significant process of analyzing and optimizing the choice of all Tx and Rx parameters, and we are not interested in assessing a single detailed architecture, but rather getting higher-level answers to basic questions about the retrievability of PBL water vapor from space using the DAR technique.