Regime-dependent evolution of water vapor, cloud liquid, and cloud base before rain onset in the southeastern Alpine forelands

Esmail Ghaemi, Andreas Kvas, Ulrich Foelsche, and Pavlos Kollias

We thank the referee for the detailed and constructive comments, which have helped us improve the manuscript. The main changes are: (i) a revised abstract and introduction with a clearer motivation, (ii) instrument accuracies and limitations in Sects. 2 and 3, (iii) a cloud-fraction analysis (new Fig. 8) and a qualified interpretation of the joint IWV–LWP evolution, (iv) a precipitation overview by category (Figs. 1c and 4), (v) a new Discussion section (Sect. 4.5), and (vi) shorter case studies, with two in the main text and three in Appendix A. Our point-by-point responses are given below. Referee comments are in black italics, and our responses are in blue.

Esmail Ghaemi, On behalf of all co-authors

RC 2:

The paper investigates the typical behavior of liquid water path (LWP), integrated water vapor (IWV; the first two derived from microwave radiometer, MWR), and cloud base height (CBH) at the Wegener Laboratory, Austria, before the onset of precipitation. This involves a statistical approach where a large number of precipitation events are identified from rain gauges and classified with the help of polarimetric X-band radar into warm/cold and convective/rest conditions. As explained in the conclusions this could lead to a new metric for evaluating forecasting models.
While the topic is generally interesting, the paper needs substantial revisions (in particular how it is written) before it can be published. My main criticism is

1.There is no clear motivation given in abstract and introduction. The conclusions are much better written and help to put the paper into context. A discussion on physical processes, precipitation efficiency, the identified gap in understanding, and the addressed spatial/temporal scales is necessary for a paper in ACP, which should address a broader audience. Because a major result is a new methodology for model evaluation, the authors should highlight that, nowadays (particularly within the ACTRIS network and operational radar), similar observations are available from multiple sites, enabling broader application of the technique developed here
Thank you for this comment. We have revised the abstract (within the length limit) and the introduction. The revised introduction now discusses the processes between cloud formation and surface precipitation, the difficulty of estimating precipitation efficiency from observations (Kukulies et al., 2024), the role of advection, and the different spatial and temporal scales of convective and non-convective precipitation.

We also state the gap addressed by this study more clearly: local column observations cannot close the full chain from moisture supply to surface precipitation, but they can show when and in which order the column moistens, liquid water builds up, and the cloud base lowers before rain reaches the surface, which models should also reproduce. Regarding the scales, we note in Sect. 2 that the 2.5 min sampling corresponds to about 1.5 km, which is comparable to the grid spacing of convection-permitting models (as suggested). Finally, as suggested, the introduction and the Discussion name ACTRIS, and the abstract mentions transferability to other sites (Illingworth et al., 2007; Laj et al., 2024), so the method can be applied at other sites and used for model evaluation.

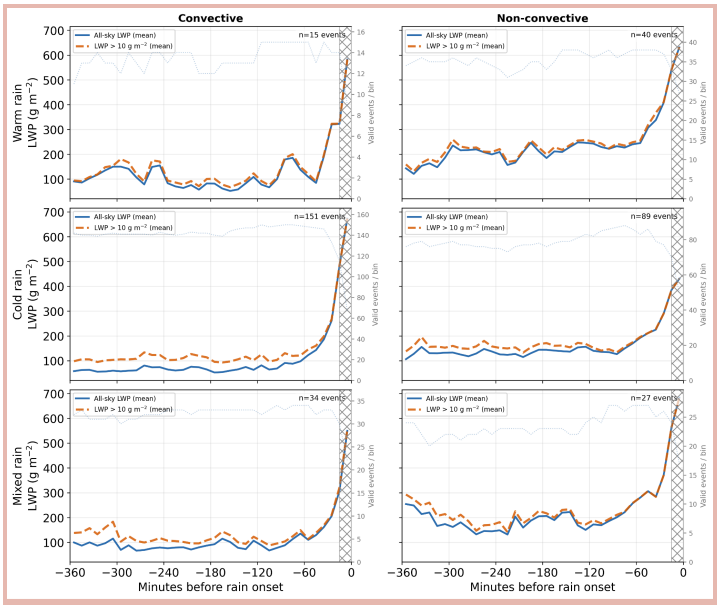
2. No information on the accuracy of the measurements is given. I point out several details about that later and suggest integrating this information into Section 2, which is rather short. It does not provide the basis for understanding the limitations of the observations (parts of it are in section 3 or even later, but it would be better to have it together for all instruments)

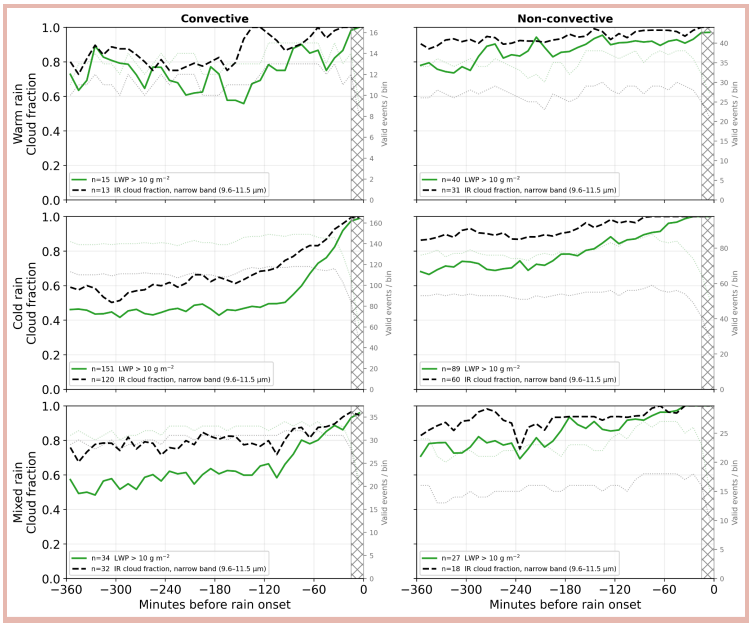
Thank you for this comment. We have added instrument-specific accuracy information to Table 1, and we describe the measurement accuracy and the limitations of the observation techniques in Sects. 2 and 3. We also discuss the vertical grid spacing and the effective vertical resolution of the radiometer temperature profiles and their effect on the −5 °C isotherm, and the differences between GNSS- and MWR-derived IWV. Further details are given in our response to the comment on Section 2 and Table 1.

3.The methodology should be sharpened: The choice for the considered averaging and sampling times needs to be made clear. There is no consideration of cloud fraction (CF); it is not even discussed. In my opinion, CF is the main reason for the finding "The joint evolution shows that convective events first move towards higher IWV and then towards a strong LWP increase near onset, while both parameters increase together in non-convective events". I suspect that conditional averaging (mean of LWP when present) would lead to a rather different results as downdrafts in convective conditions.

We thank the referee for raising this concern. To check the effect of removing cloud-free samples and of the temporal averaging, we compared all-sky LWP with LWP conditioned on values of above 10 g m^{-2}, using the original higher-frequency measurements and averaging intervals from 30 s to 10 min. The pre-onset behavior remained mainly similar, including the relatively low LWP well before convective onset (figure below, left).
We agree with the referee that cloud occurrence contributes to the convective signal. In the revised manuscript we therefore (i) state explicitly in Sect. 3 that the LWP composites are conditional averages over cloudy time steps (LWP > 10 g m^{-2}), (ii) show the cloud fraction in a new figure (figure below, right; Fig. 8 in the manuscript), placed directly before the joint IWV–LWP analysis, and (iii) qualify the interpretation of the joint evolution: in convective events, part of the late LWP increase reflects the increasing occurrence of cloud over the radiometer, i.e. the arrival of the convective cloud, and not only condensation within an existing cloud. For all convective events together, about one third of the increase in all-sky mean LWP between −120 and −15 min comes from the increase in cloud fraction (0.52 to 0.97), and about two thirds

from the higher mean LWP of the cloudy time steps. This qualification is now also included in the abstract and the conclusions. Regarding the averaging and sampling times, Sect. 2 now explains that the 2.5 min interval corresponds to the radar volume-scan duration and, for an advection speed of 10 m s^{-1}, to about 1.5 km, the choice of the 6 h window is explained in Sect. 3.





4.It is assumed that the four-year statistics are large enough statistics, and all events are weighted equally. Similar to a previous study for this site, they also look at the maximum precipitation above the 90th percentile but don't consider other metrics. I think the radar data includes much more information to characterize the precipitation events, for example, on the precipitation pattern (are the ground-based measurements in the center or at the edge) and propagation direction and speed.

We appreciate this point. Radar observations contain information about the position, movement, and spatial structure of each precipitation system. However, we have already separated the sample into six warm-season rain-type/regime categories, and some of these categories contain relatively few events. A further subdivision would reduce the sample sizes and make the composite statistics less stable. We therefore keep the present classification, while recognizing that the storm position and advection relative to the radiometer can affect the observed pre-onset signatures. We plan to investigate these effects in future work using a larger radar data set (e.g. operational C-band radars and other radars).

We have added this limitation to the Discussion (Sect. 4.5). We would also like to clarify that our event selection is not based on a precipitation percentile: every event that exceeds 0.1 mm h^{-1} for at least 15 min at the radiometer site is used. The 90th percentile (P90) in our method refers to the spatial distribution of the echo-top heights around the radiometer. Each event contributes equally to the composites (now stated in Sect. 3), and the number of valid events in each time

bin is shown on the right axis of all composite figures, so the reader can judge the robustness of each category.

5.Finally, I find it rather strange that the case studies are presented after the statistical summary. I would first introduce the methodology with one or two case studies and then move to the statistics as this is the main result to be discussed in the final conclusion section. Also, the discussion of the cases is rather descriptive and lengthy, tiring the reader so maybe even more can be moved into the appendix.

We appreciate the referee's comment. We have shortened the descriptive and lengthy parts of the case-study section and moved the warm convective case (19 August 2024) to Appendix A. We considered presenting the case studies before the statistics. However, the cases are selected and interpreted with respect to the composite behavior of each category, so they are easier to follow after the composites are introduced. To keep the structure of the manuscript close to the original, we therefore keep the order, but we have added a closing paragraph that compares the two cases with the composites, and added a short discussion after the case studies, so that the transition to the conclusions is no longer abrupt.

DETAILED COMMENTS

Abstract: The first sentence is rather detailed
We have rewritten the beginning of the abstract, which now starts with the motivation.

L23: “these parameters”. What does this refer to? In the sentence before you talk about processes and scales. As noted in my major concern 1, I would expect more discussion of the physical processes and the difficulty of determining precipitation efficiency from observations (see for example Kukulies et al., 2024, https://doi.org/10.1029/2024JD041924).
We have clarified the wording in the introduction: "these parameters" has been removed, and LWP and IWV are now introduced explicitly as two column-integrated parameters. We also added a short discussion of why estimating precipitation efficiency from local observations is difficult. Our measurements describe the evolution of column moisture, liquid water, and cloud-base height before surface rain onset; they do not directly measure the conversion efficiency of cloud condensate into surface precipitation. We have clarified this limitation in the manuscript.

Starting L33: Rather detailed results from previous studies are given but the main challenge of the study is ignored: How can local measurements contribute to the understanding of the end result of a complex chain (surface precipitation) that is influenced by atmospheric dynamics on various scales. The authors need to think a bit more broadly in the introduction as for example advection is never mentioned.
We agree that column measurements and surface precipitation reflect different scales. We have added advection and storm position as possible sources of mismatch, and clarified that these composites describe observed pre-onset signatures, not the full chain of precipitation formation.

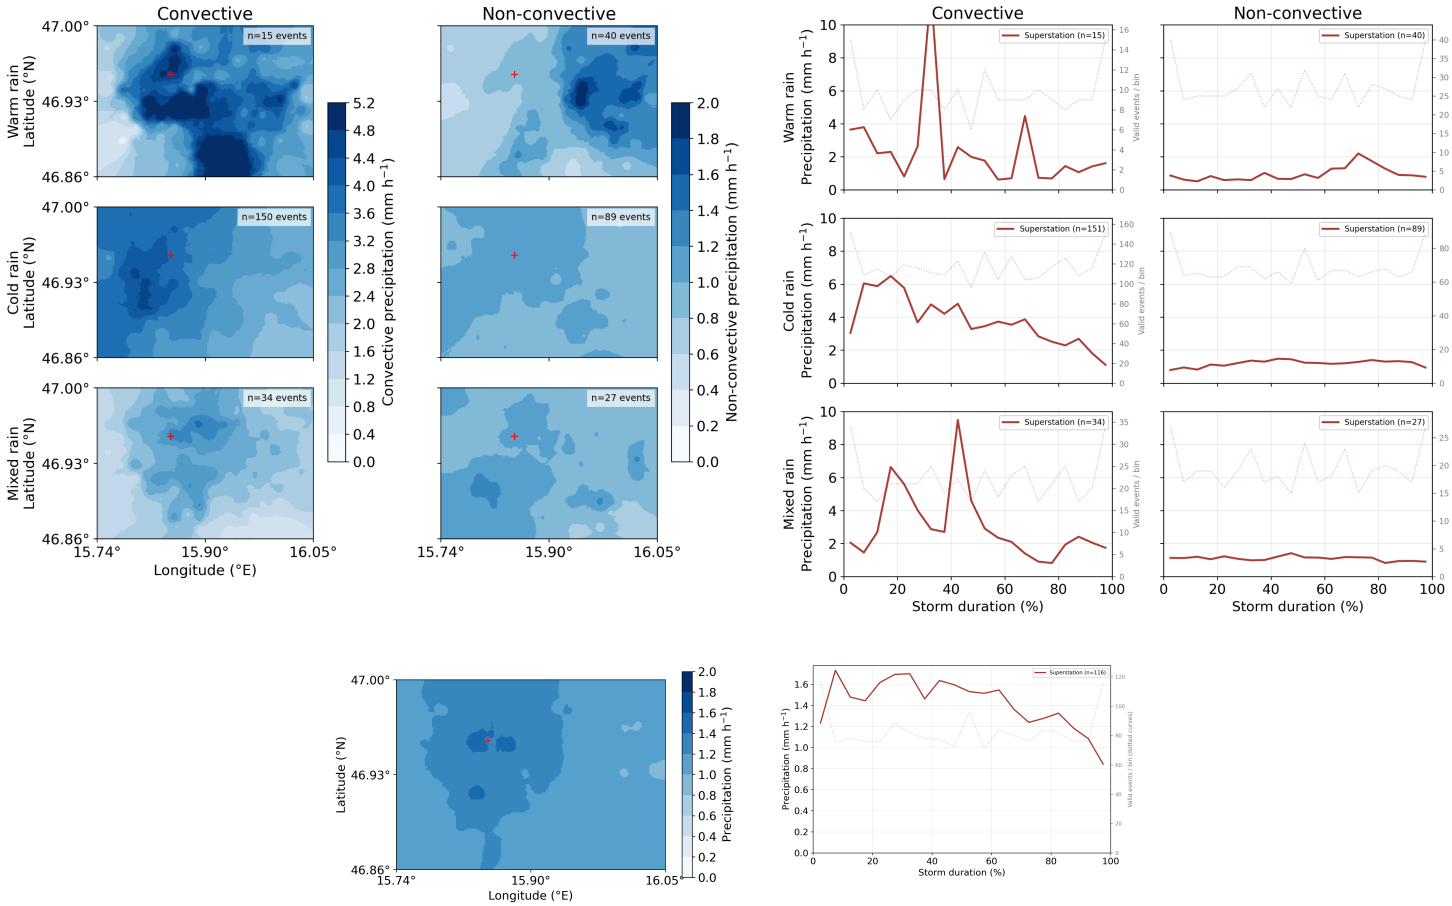
In the introduction, we now also state what local column measurements can contribute to this chain, namely the timing and the order of moistening, liquid-water build-up and cloud-base lowering before rain reaches the surface. We also note that quantifying the advective contribution would require horizontally resolved wind and moisture observations, for example from a network of profiling sites (Wagner et al., 2022).

L45: The distinction between convective and non-convective scenes is a good approach to move toward more local - less advective - dominated condition and can be strengthened here. The rather detailed description of thresholding should be replaced by the idea behind and not the details.

We appreciate this suggestion. We have strengthened the introduction to explain why the convective/non-convective distinction is useful and have kept the technical classification details in the methodology.

L 51: This question is not explicitly answered in the conclusions
We have revised the second question, which now refers explicitly to the definition used in this study (events with radar echo tops below or above the −5 °C isotherm). The Conclusions are now organized along the three questions, and each question is answered explicitly.

L62 onward: Refer to Fig.1 How large is the area? What is the typical precipitation? There is one sentence referring to the VB – but how much does this contribute? How does this relate to the investigations? Fig.1a shows the impressive ground-based station network. As far as I understand, these sites all measure precipitation. Why don't the authors show the average precipitation rate over all events studied and separated into their different categories (warm, cold, convective..) such that the reader gets an overview how different their distributions are. In fact, it might also be promising to distinguish the events based on the circulation weather type, as these lead to different interactions with orography.
We appreciate this comment. We have added the horizontal extent of the gauge-covered study area (22 km × 16 km) and a precipitation overview based on the gridded station-network product and the local gauge measurements, separated by regime and rain type (Figs. 1c and 4, figures below). We also added a short description of the precipitation characteristics of these categories. The topography is moderate, and our earlier regional analysis suggested a relatively small orographic influence on precipitation within the study area (Ghaemi et al., 2021). The contributions of Vb cyclones and of the circulation weather types were not quantified separately, as such an analysis needs a larger data set with more spatial coverage.
We agree that a separation by circulation weather type is a promising extension and mention it as future work in the Discussion (Sect. 4.5).



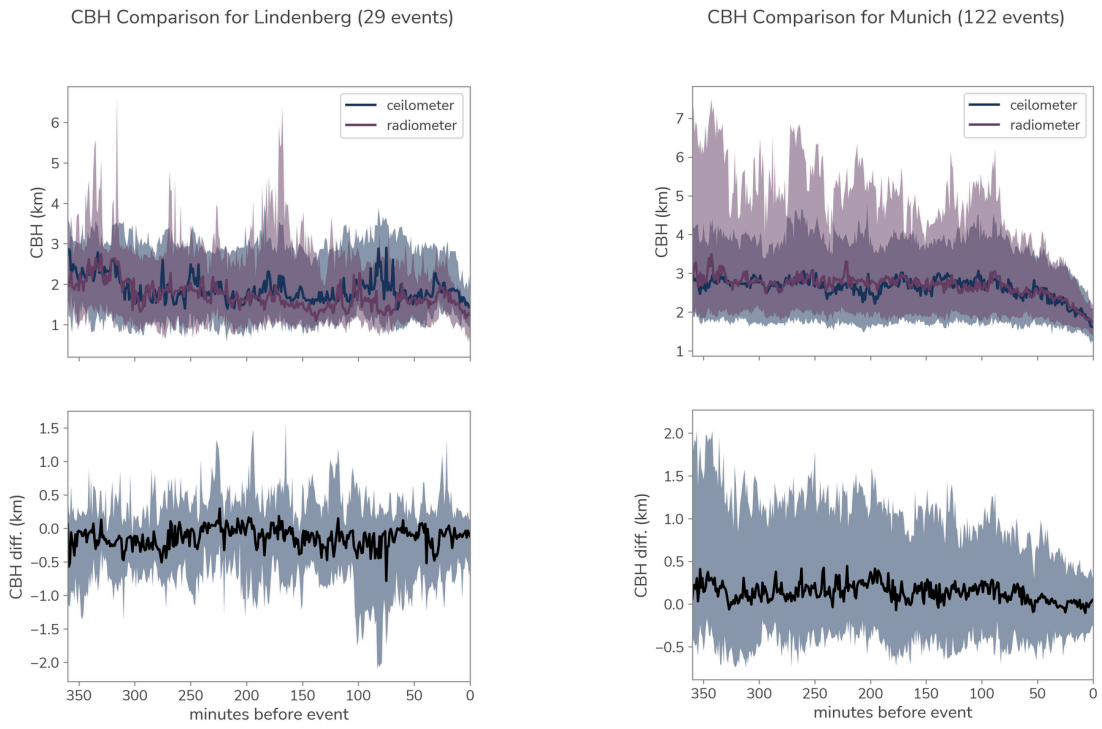
L79: How far are the precipitation measurements away from the central = radiometer (define) site?

The instruments at the central Raabtal site are collocated within a few meters at the superstation. For the event detection, we use the average of four rain gauges: the gauge at the superstation and three neighboring gauges at distances of 225 m, 370 m, and 835 m from the superstation. We have added this information (mean distance of about 500 m) to Sect. 2.

Section 2: Table 1, which gives an overview of the data, is not referred to. There is no information on the key products, their uncertainty, and their temporal sampling in the text. For all data, a common resolution of 2.5 min is used. What is the reason for that? Typically, MWRs measure with second resolution, while GNSS have longer integration times as several slant path measurements to satellites are needed. The difference in spatial (and temporal) samples causes a different sensitivity to water vapor variability, especially in convective situations (Steinke et al., 2015, doi:10.5194/acp-15-2675-2015). Thus, this partly explains IWV differences observed between both instruments, as for example stated on line 236. In general, it was not clear when you use MWR or GNSS and why? Another missing information is how CBH is derived. The IR temperature is affected by water vapor emission between cloud base and surface so there is a sensitivity to IWV here. Also, I don't know how meaningful the measurement is during virga precipitation. Another point, it is the temperature profile: Section 2

should explain that the MWR retrieval is used (always?) and what the uncertainty/limitation is. In summary, there is important information to be provided here.

Thank you for this comment. Table 1 is now referred to in Sect. 2, and it gives the accuracy and the temporal sampling of each product.
The common resolution of 2.5 min is chosen to align the radiometer/GNSS time series with the duration of a radar volume scan. The radiometer data are aggregated from their second-resolution retrievals, while the GNSS observations are constrained in time to increase redundancy. When comparing periodograms from both measurement techniques, we can see that the effective temporal resolution is very similar, i.e. GNSS is not overly smooth compared to the MWR. The GNSS IWV retrievals are based on multi-GNSS data, which provide an improved sampling of the atmosphere. This at least partly mitigates the sensitivity issue. Neither effect can explain the inter-technique bias of 1 kg m^{-2}, which is persistent across seasons and in both rain and no-rain conditions.
We agree that the larger GNSS footprint makes it less sensitive to small-scale water vapor variability, in particular in convective situations (Steinke et al., 2015). We added this reference and a short discussion to Sects. 3 and 4.2, where we also state that the footprint mainly affects the short-term variability, whereas the persistent 1 kg m^{-2} offset cannot be explained by it.
Regarding the use of MWR and GNSS: MWR IWV is used together with LWP, because both are retrieved from the same measurements. GNSS IWV is shown in addition as an independent estimate that is not affected by the wet radome, and it is therefore also available in the last minutes before onset. This is now explained in Sect. 3.
CBH is derived by matching the infrared brightness temperature with MWR-derived temperature profiles. As the referee rightly states, IWV affects the brightness temperature, with higher IWV effectively lowering the estimated CBH. To reduce the IWV impact, we apply the dual-angle method of Smith and Toumi (2008). We further evaluated the method at two ACTRIS Cloudnet sites with similar climatological conditions (Lindenberg and Munich), where comparable instrumentation and independent CBH information (ceilometer) are available (figures below).



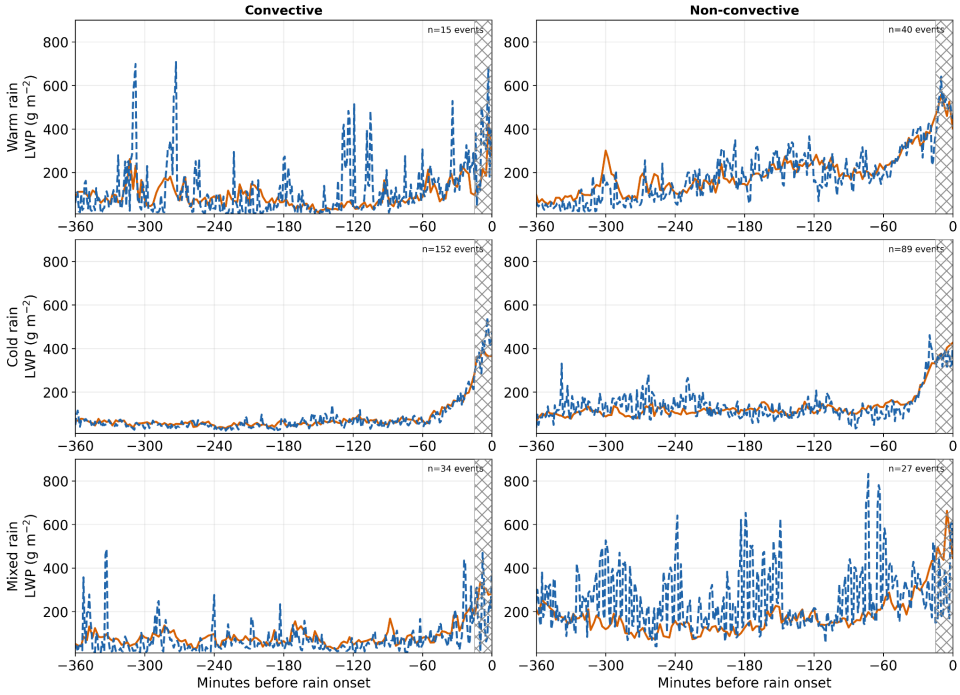
The comparison shows that the method is consistent with the ceilometer observations close to the event onset, with larger variations at the beginning of the 6 hour window. These larger variations may be attributed to thinner and higher clouds, where the drawbacks of the IR/MWR combination are amplified.
The infrared radiometer detects the lowest radiatively opaque level. During virga, or in the last minutes before onset, precipitation below cloud base may also affect the retrieved CBH, so the retrieved CBH may be lower than the true cloud base in the last minutes before onset. The ceilometer has a similar limitation. This effect cannot be separated with our instruments and is one reason why the last 15 min are marked as uncertain. A sentence is added on this in Sect. 3.
Regarding the temperature profile: the temperature profiles are taken from the microwave radiometer retrieval. They are provided on a vertical grid with a spacing from about 20 m near the surface to about 100 m at the typical height of the −5 °C isotherm (Kvas et al., 2026), and this grid spacing is shown as gray shading in the case-study figures. However, the effective vertical resolution of the retrieved profiles is coarser and decreases with altitude, due to the limited vertical information content of passive microwave observations (Crewell and Löhnert, 2007). This is now stated in Sect. 3.
We amended the manuscript with uncertainty information on the retrievals and a brief discussion of the MWR retrieval limitations and their impact on the CBH estimation.

My major concern with the 2.5 min resolution for LWP is that the information on broken cloud situations (typical for convective conditions) is lost. Are the original data still available? At least it would be good to have the information on cloud fraction in this time period. LWP is likely lower in convective than in stratiform cases if it is averaged over cloud free periods. Maybe the

authors can argue that 2.5 min could mimic an atmospheric model grid cell of 1.5 km, assuming cloud advection of 10 m/s.

We appreciate the referee's comment, and we agree that a 2.5 min resolution may be coarse, especially in convective situations. However, using a finer temporal resolution (e.g. sub-minute) gives noisy results. Following the referee's suggestion, we have added a sentence to the manuscript explaining the choice of the temporal resolution. We adopted the referee's argument in Sect. 2: for an advection speed of 10 m s^{-1}, 2.5 min corresponds to about 1.5 km, which is comparable to the grid spacing of convection-permitting models.

The original data are still available, the figure below shows the composites based on the sub-minute data (dotted blue lines indicate sub-minute data).



L 99: "which are not derived from interpolation" what does that mean?

Here, it simply means that only samples not flagged by quality control are used. We have removed this part.

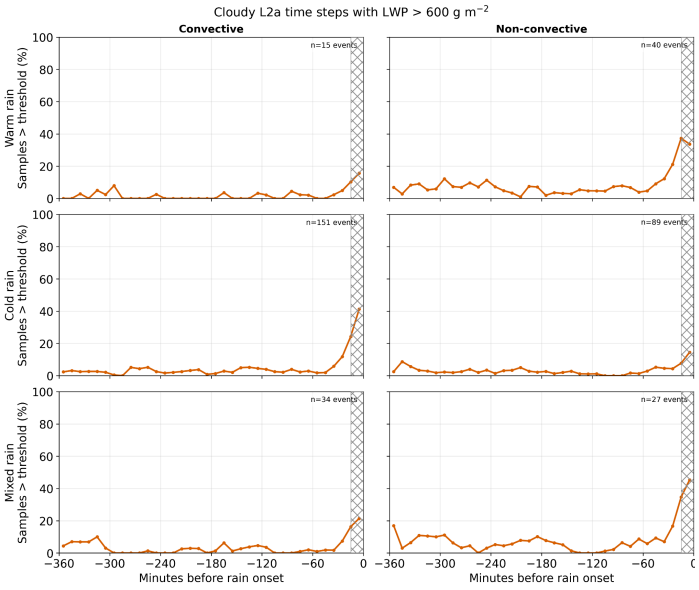
L105: The multiplication is trivial

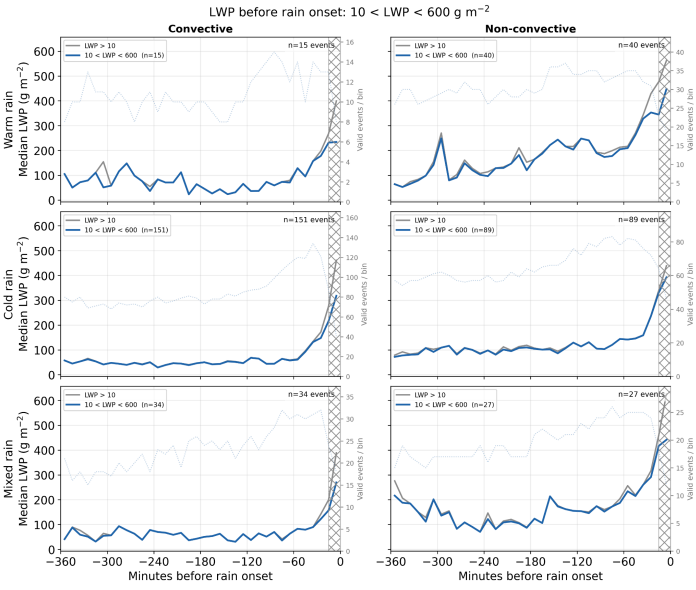
We have removed this explanation.

L 114: It was unclear there the temperatures come from. BTW, did you perform a sensitivity study especially because the retrieval is most uncertain during thick clouds?

The temperature profiles are retrieved from the microwave radiometer brightness temperatures in the oxygen absorption band, with the same retrieval framework as LWP and IWV (Böck et al., 2025). The −5 °C isotherm (and the freezing level used in the convective scheme) is derived

from these profiles. Because the retrieval is less reliable during rain and with a wet radome, we average the isotherm height between 60 and 30 min before onset (or over the 2 h before onset if all time steps in this window are flagged). We have clarified this in Sect. 3. To test the sensitivity of the rain-type classification to the temperature retrieval, we shifted the −5 °C isotherm height of each of the 356 warm-season events by ±500 m, which is comparable to the effective vertical resolution of microwave radiometer temperature profiles in the lower troposphere (Crewell and Löhnert, 2007). With ±500 m, about 19 % of the events change type, almost entirely at the boundaries of the mixed type: 54 % of the mixed events become warm rain with +500 m and 66 % become cold rain with −500 m, and 47 % of the warm events become mixed with −500 m. Cold-rain events are more robust (14 % become mixed with +500 m and none change with −500 m), and only 3 of 55 warm events change directly to cold rain. Therefore, the separation between warm and cold rain is robust, while the mixed type is more sensitive to the uncertainty of the isotherm height, as expected from its definition. A short summary of this test is now given in Sect. 3.
Additionally, we analyzed all radiometer data to check which fraction of LWP values exceeds 600 g m⁻², above which the retrieval uncertainty increases as shown below. Before −30 min, these samples account for less than 6.3 % of the selected samples within each category, although individual time bins reach up to 20.7 %. Excluding them has a negligible effect on the median LWP before about −60 min; the largest differences occur within the last 30 min before onset, which we already interpret with caution.





L135: You don't mention what rho_hv is? – why do you have 0.85 here and 0.9 later?
The three thresholds have different purposes: echoes with ρhv < 0.70 are removed as non-meteorological; for quasi-stationary echoes (|vrad| < 0.5 m s⁻¹), a stricter threshold of 0.90

is used to remove ground clutter, and the convective criteria are evaluated only where ρhv ≥ 0.85, so that mixed or uncertain targets do not contribute to the convective score. Each threshold is now given once, together with its purpose, in Sect. 3 and in Fig. 3.

L136: The polarimetric variables are not defined and explained. I think it is a good idea to go not too much into detail here but maybe you can put the threshold values into the table with the sensitivity study in the appendix and only explain the physical principle in section 3.
We have added concise physical definitions of Zh, ZDR, KDP, phv, radial velocity, and spectrum width in Sect. 3. The detailed threshold and weighting choices are summarized in Fig. 3 and Table A1 to keep the main text compact.

Section 4.1: The monthly mean IWV/LWP analysis is a bit separated from the rest of the study. Nevertheless, I find it valuable but it needs to be put into context, i.e.. saying how unique this is, as observed LWP climatologies only exist over ocean (MAC-LWP; MODIS) and that large discrepancies exist between them (Lohmann and Neubauer, 2018 https://doi.org/10.5194/acp-18-8807-2018)

We retain the monthly context because it provides the climatological background for the regime-dependent composites. We have added one sentence noting that long ground-based LWP records over land are available only at a limited number of sites, while satellite LWP climatologies such as MAC-LWP are restricted to the oceans (Elsaesser et al., 2017), and that large discrepancies exist between different LWP climatologies (Lohmann and Neubauer, 2018).

L 166: How many samples are necessary to get a robust monthly estimate? Do you apply any thresholding? I ask because in Fig.4, March 2022 has a monthly mean of 500 gm-2. As this is a value often connected already with precip (and retrievals of MWR are not made for that), I doubt that this is a representative monthly mean. As MWR estimates exclude precip events, the data are biased. This also affects the intercomparison with ERA5 and might explain the higher values there. For IWV this bias can be estimated from the GNSS data – it would be interesting to see if these data give the same results on the "precip bias" as found by Henken et al. 2020, https://doi.org/10.3390/rs12071170

We appreciate the referee's comment. The minimum number of samples per month in the plot was 40 time steps for the cold-season events, while for the warm-season events there were at least 250 time steps. Regarding the high value in March 2022, this is a good catch. An event on 1 April 2022 started at 04:50 UTC, so its 6 h pre-onset window extended into March 2022, and our algorithm assigned these time steps to March, while classifying the event as a warm-season event. This single borderline event caused the high monthly value. We have now assigned all time steps of this event to April 2022. We now also state in Sect. 4.1 that the MWR monthly statistics exclude precipitating periods, so the comparison with ERA5 is affected by this sampling difference. Regarding the GNSS-based estimate of this precipitation bias (Henken et al., 2020), we agree that this comparison is interesting and leave it for future work.

L167: Repetition of L99 – still don't understand "interpolated"

We have removed the term "interpolated" (see our response to the comment on L99).

L170: How are warm and cold season defined?

We define the warm season as April to September, the period in which most convective events occur in the region (e.g., Kirchengast et al., 2014, Schroeer and Kirchengaast 2018, Kvas et. al., 2026), and the cold season as October to March. This is stated in Sect. 3.

L175: Hersbach, H.L179: Figure captions don't belong into the main text. The text should be understandable without looking at figures. Figures are there to support your statements. This comment also applies to other parts of the paper,

We have corrected the reference formatting and rewritten the passages that repeated figure captions, focusing instead on the interpretation of the results.

L186: What does "changes" mean? Differences between both data set? LWP variability?

Here we meant the differences between the two data sets. We have rephrased the sentence accordingly.

L188: As the scale issue is of high importance for the study the choice of the considered interval should be motivated earlier.

We agree with the referee's comment and have moved this part to the methodology (Sect. 3).

L193: I wouldn't call it a jump, but more an exponential decay. The authors might consider to reduce bin size in Fig.5 and similar ones (or increase bin size with increasing time) as the grey bars dominate the plot. This may even allow adding a smoothed line / fit to become more quantitative.

We now use "rapid increase" rather than "jump". We tested the visual presentation, but kept the confidence bands, because they show the sampling uncertainty.

L202: Can you better quantify the impact of the depth of convection using radar rather than distinguishing between warm=shallow and deep=ice. In fact that you have this definition should be made clear in the second question posed in the introduction.

We tried to estimate the physical cloud thickness before rain from the radar echo-top height. However, this resulted in even fewer cases and therefore noisier results, simply because many pre-event clouds could not be detected by the X-band radar.

This is consistent with the limited sensitivity of the X-band radar to non-precipitating clouds. As suggested, the second question in the introduction now states explicitly that the depth is defined by the echo top relative to the −5 °C isotherm.

L222 Did I miss a figure on the diurnal cycle? If there is none it might be good to better quantify the effect, e.g. by giving the median times of day when the events happen in the different classes.

We have added the distribution of the pre-onset observation times for each category to the appendix (Fig. A4).

L224: What is the "moist" part of the day? It strongly depends which moisture quantity you are looking at. In IWV, the diurnal cycle mostly can't be detected.

We agree that "the moist part of the day" is imprecise. We have removed this expression and now refer to the hours in which our convective events occur more frequently.

We also agree that the diurnal cycle of IWV is weak. Therefore, we no longer attribute the higher IWV in convective events to the time of day. It is more likely related to the fact that convective events occur mainly in the warmest months and in warm air masses. The sentence in Sect. 4.2 has been revised accordingly.

L227: In general, the discussion of the rather different ranges of IWV and LWP needs to be clarified at the beginning. There are many descriptions along the line that "IWV increases less than LWP". The statement might change if relative units are used as for example a 10 % increase in IWV (e.g. 20 to 22mm) and an 10% LWP increase (eg,. 0.10 to 0.11mm) are compared.

We agree that absolute ΔIWV and ΔLWP values should not be interpreted as a comparison of relative rates.

We have added a sentence at the beginning of the joint-evolution discussion (Sect. 4.2) noting that IWV and LWP differ by about two orders of magnitude, so the same absolute change corresponds to very different relative changes (an IWV increase of 3 kg m^{-2} is about 10 % of the column vapor, while an LWP increase of 0.3 kg m^{-2} can be several times the pre-onset LWP).

L228 "more ice particle column (cold rain) holds less water vapor. That sounds rather strange. The highest IWV worldwide is found in deep (ice) convection in the tropics. In mid latitudes, I think it mostly depends on how high the tropopause level is…and tropopause is mostly higher in warm airmasses.

We agree that the earlier wording incorrectly implied that ice-rich clouds necessarily contain less water vapor. We have removed that causal explanation. The lower IWV in our cold-rain subset is an empirical class difference, likely affected by the sampled air-mass temperature and synoptic conditions.

L235: onward – see discussion on the different footprint of MWR and GNSS.

We agree. We have added to Sect. 3 that GNSS IWV has a larger spatial footprint than the MWR, as several slant delays are mapped to zenith, and that this makes GNSS less sensitive to localized moisture gradients, in particular in convective conditions (Steinke et al., 2015). In Sect. 4.2, we note that the footprint difference mainly affects the short-term variability, while the persistent offset of about 1 kg m^{-2} between the two instruments is found across seasons and also in no-rain conditions, so it cannot be explained by the footprint alone (see also our response to the comment on Section 2).

L242: what is a "milder" CBH?

We meant that the decrease in CBH is more gradual and less pronounced, and have revised the wording accordingly.

L247: more persistent clouds? Do you have information on cloud fraction?

Yes. Based on the cloud fraction from the infrared radiometer (available for fewer events), the cloud fraction before onset is higher in the cold season (not shown), and in non-convective events than in convective events (Fig. 8).

L249: Very sloppy text – describe what you investigate

We have rewritten this paragraph. It now first describes what is investigated (the joint evolution of the IWV and LWP anomalies, Fig. 9) and then gives the interpretation.

L258: Are the correlations significant – for r=0.4 the explained variance is only 16%

We agree with the comment. Both correlations are statistically significant ($r = 0.46$, $n = 136$ and $r = 0.39$, $n = 123$; $p < 0.001$), but they explain only about 21 % and 15 % of the variance. We now give the significance and the explained variance in Sect. 4.2 and describe the relation as a weak-to-moderate tendency rather than a strong dependence.

L270: Was this a cold front passing? For the case studies you should be able to check the synoptic situation.

This case was local convection, and the temperature drop was caused by the passage of a convective cold pool rather than a cold front. The weather charts for this day (by GeoSphere Austria) show no frontal passage over the region at this time. The rapid temperature decrease together with the increase in wind gusts is consistent with a cold pool. This is now stated in the manuscript.

L274: You mention which data you use in Section 2 so you don't need to repeat Level2… radar derived precipitation is sufficient. The sentences is again just a figure caption which also makes the section tiring to read.
We have shortened the introductory figure descriptions in the case studies.

L280: Before rain starts – it might also be mismatching and shear not seen by the radar. How far down does the radr get without being affected by clutter?
We agree that a spatial mismatch can contribute in this event, as the wind was strong. We have added beam height, vertical shear, advection, and sub-cloud evaporation as possible reasons for the timing difference between radar echoes aloft and rain at the surface (Sect. 4.3). Within a range of about 30 km, the radar data are not significantly affected by clutter.
Over the radiometer site, at 14.5 km from the radar, the lowest elevation angle of 1° corresponds to a beam-center height of about 270 m above the radar altitude (about 640 m). Since the radiometer site is at about 320 m, the lowest beam is about 590 m above the ground at the radiometer site.

L300: was this a single cell storm?
We consider this event as an organized convective precipitation band rather than a single cell, potentially associated with a larger-scale air-mass boundary. However, since only a few gauges captured rain, and only for a short period, partial sub-cloud evaporation may have reduced the rainfall reaching the gauges.

Fig. 11: Why do you have so much ZDR up into high altitudes before the event?

At the time when these spurious ZDR values occur, the storm passed through the region where the radar switches from short-pulse to long-pulse mode. We sometimes observe ZDR/KDP jumps there, which are not always captured by the quality control.

L308: 0.6 mm/h is rather weak – is this typical for the non-convective class? In fact it would be nice to see how the precipitation rate frequency of occurrence differs between the different categories.
We have added the category-dependent precipitation intensities in Fig. 4, which place the ~0.6 mm h^{-1} case in context.

It would be nice to see how the different cases fit into the overall statistics – how are they chosen. The transition from the end of 4.3 to the conclusions is rather abrupt. There is no discussion about what was learned from them. In my major comment 5 I suggest to move them before the statistics.
The cases were selected to illustrate the contrasting pre-onset behaviors, and because radar and radiometer data are complete in the six hours before onset, not as statistically representative examples. We have added a short paragraph comparing their IWV/LWP/CBH signatures with each other and with the composites, and we have moved additional cases to the appendix. The case studies are now followed by the Discussion, which links the case studies and the composites to the conclusions. Regarding the order, please see our response to major comment 5.
L333: How do you know that where the liquid is in the cloud? It might well be that there is a lot of supercooled liquid transported aloft by strong updrafts but can you proof it?

Reply to Referee 2 (RC2)

We cannot infer the vertical location of liquid condensate from column-integrated LWP. We now describe the ZDR column as consistent with oblate hydrometeors above the melting level, while avoiding a direct claim about the distribution or phase of all condensate.

L380: Be careful: This is different in relative units. However, here one has to be careful as well because LWP is bounded by zero while IWV is continuous.

We agree. We now refer to their temporal sequence and distinct mass scales, with the caution that relative LWP changes are unstable when the baseline is near zero.

We keep one absolute comparison (the LWP increase of about 0.4 kg m^{-2} versus the IWV increase of about 3 kg m^{-2}), because it refers to the column water mass budget, and we now state this explicitly ("in absolute (mass) terms", Sects. 4.2 and 5).