General Response to the Editor and Reviewers

We would like to again sincerely thank the Editor and both reviewers for their careful, detailed, and constructive evaluations of our manuscript. The reviewers raised a number of important and highly relevant comments concerning instrument characterization, methodological description, the definition and interpretation of physical constraints, the treatment of platform-attitude data, and the interpretation and scope of the retrieval results. We fully recognize that these comments are highly valuable for improving the scientific rigor, technical completeness, and clarity of the manuscript.

We also fully acknowledge that the original version of the manuscript contained several aspects that were insufficiently clarified, not rigorously justified, or not cautiously delimited. In particular, the main concerns identified by the reviewers include: (1) the description of the microwave radiometer performance and uncertainty was somewhat generalized; (2) the physical meaning of the retrieval constraints, especially the so-called "lapse-rate constraint," was not sufficiently clarified; (3) the treatment of attitude data relative to the radiometer integration time was not explicitly described; (4) the interpretation of the oscillatory behaviour in the retrieved profiles was not sufficiently rigorous; (5) a systematic quantitative evaluation of the improvement associated with incorporating attitude information was lacking; and (6) no direct comparison with the built-in radiometer retrieval algorithm was provided in the original manuscript.

In response to these concerns, we have carried out a substantial and systematic revision of the manuscript. All modifications have been clearly marked in the revised manuscript, and corresponding page and line numbers are provided in the detailed responses below. The main revisions can be summarized as follows.

1. We strengthened the description of the instrument background and clarified the boundary of the uncertainty discussion. In the revised manuscript, we now explicitly distinguish between the typical performance ranges reported for comparable classes of ground-based microwave radiometers and the instrument-specific technical specifications of the QFW-6000 used in this study. We also added Table 1 to summarize the key technical specifications of the instrument and explicitly state that these values provide technical background information rather than a complete field-based uncertainty characterization under both fixed-site and buoy-based deployment conditions. In addition, we now clearly acknowledge that the present work does not include a dedicated comparison experiment between stationary land-based deployment and buoy-based deployment, and therefore

should not be interpreted as a complete metrological characterization of platform-dependent instrument uncertainty.

2. We clarified the mathematical formulation and physical meaning of the retrieval constraints. In particular, we now explicitly state that the so-called “lapse-rate constraint” is not a strict monotonic decrease constraint, but is implemented as an adjacent-layer continuity constraint. We also provided the specific values of the corresponding constraint parameters. Furthermore, in order to respond directly to the reviewers’ concern that these constraints might artificially suppress strong inversions or local gradient-transition structures, we added a representative case study (Fig. 10) to illustrate the actual effect of the imposed bounds and adjacent-layer continuity constraints on the retrieved profiles.
3. We supplemented the manuscript with a clearer description of the temporal treatment of attitude data. The revised manuscript now explains that the attitude sensor samples at a higher frequency than the microwave radiometer and that all attitude samples within each radiometer integration window are synchronized and averaged to obtain an effective attitude state. Based on this averaged attitude state, an effective zenith angle is calculated and then used in the forward radiative transfer model, thereby ensuring consistency between the time-integrated radiometric observation and the observation geometry used in the retrieval.
4. We revised the interpretation of the results, especially regarding the oscillatory behaviour of the mean retrieved temperature profile. In the revised manuscript, we no longer describe this behaviour as a “typical feature” of microwave radiometer retrievals. Instead, we interpret it more cautiously as the combined result of reduced information content in some altitude ranges, the discrete layer-by-layer optimization framework, and the current regularization setting. We now also explicitly acknowledge that the present regularization strength may not be fully sufficient to suppress small-scale oscillations under weak observational constraints.
5. We added a systematic quantitative analysis of the impact of attitude correction. This includes: (1) grouped statistics under different attitude-angle ranges; (2) a comparison between retrievals with and without attitude correction; and (3) a theoretical sensitivity analysis of brightness temperature to viewing angle (Fig. 8).
6. Following the reviewers’ recommendation, we added a direct comparison with the built-in radiometer retrieval algorithm (Table 5), in order to better demonstrate the practical added value of the proposed method.

7. We tightened the scope of our conclusions. The revised manuscript now states more explicitly that the reported results are based on 38 matched radiosonde cases obtained during a short-term nearshore campaign in Jiaozhou Bay. Therefore, the present results should be interpreted primarily as evidence of retrieval feasibility and systematic-error characterization under the tested nearshore buoy-based conditions, rather than as a generalized validation for all marine environments or open-ocean applications.

Overall, we believe that these revisions have substantially improved the manuscript in terms of methodological transparency, physical consistency, and the interpretation of the results. We again sincerely thank the reviewers for their valuable comments, which have helped us significantly improve both the scientific content and the presentation of the manuscript.

Response to RC1

L118: is that true for the QFW-6000? I still don't get why the instrument was never properly characterized? Or at least deployed at a non-moving site to see the difference between land and buoy-based deployment. The characterization of the instrument, especially the uncertainties, is not sufficient for me. The typical errors given all refer to other MWRs. The reference given, Zhang et al. (2025), describes an MWR with seven channels per antenna. Here there are eight.

Response:

We sincerely thank the reviewer for this important and highly pertinent comment. We fully agree that the original manuscript did not sufficiently clarify the characterization of the QFW-6000 instrument itself, especially with respect to uncertainty-related information. We also agree that the previous wording could easily give the impression that typical performance values reported in the literature for other microwave radiometers could be directly applied to the specific instrument used in this study. In retrospect, we recognize that this was not rigorous enough and could lead to misunderstanding regarding the actual level of instrument-specific characterization provided in the manuscript.

To address this issue, we have made a series of coordinated revisions in the revised manuscript.

First, in Section 2.1 of the revised manuscript (Page 5, Lines 117–134), we rewrote the relevant paragraphs to explicitly distinguish between two different types of information: (i) typical uncertainty and performance ranges reported in the literature for comparable classes of ground-based microwave radiometers, and (ii) the technical specifications of the specific QFW-6000 system used in this study. The literature-based values are now clearly identified as background information for comparable instrument classes only, and no longer presented in a way that could be interpreted as a complete field-based uncertainty characterization of the present instrument.

Second, in order to provide clearer instrument-specific background information, we added Table 1 in the revised manuscript. This table summarizes the main performance specifications of the QFW-6000 used in this study, including brightness temperature sensitivity, brightness temperature measurement error, and profile-level RMSE specifications for temperature and relative humidity. We would like to emphasize that these values are extracted from the instrument's functional/specification documentation and are intended to provide engineering-level background information on the present system. At the same time, we explicitly state in the table note that these values should not be interpreted as a complete field-based uncertainty characterization under both fixed-site and buoy-based deployment conditions.

Third, we appreciate the reviewer’s observation that the system described in Zhang et al. (2025) is not identical to the one used in our study. In the revised manuscript, we now explicitly clarify that the present QFW-6000 and the one described by Zhang et al. (2025) belong to the same instrument family, but their channel configurations are not identical. Specifically, the system used in this study includes eight channels in each of the K- and V-bands, whereas the instrument reported by Zhang et al. (2025) uses seven channels per band. We also clarify that the citation of Zhang et al. (2025) is now used only as background information on a related instrument of the same family, rather than as a one-to-one characterization reference for the exact instrument used here. Although the additional channels in the present system may provide denser spectral sampling and potentially richer radiometric information from an engineering perspective, this does not automatically imply lower uncertainty and cannot substitute for dedicated instrument-specific characterization.

The revised text is as follows: “*The QFW-6000 microwave radiometric profiler used in this study belongs to the class of ground-based multi-channel microwave radiometers for atmospheric profiling.*

……

the main technical specifications relevant to retrieval uncertainty are summarized in Table 1.

Table 1. Main performance specifications of the QFW-6000 microwave radiometer used in this study

Metric	*Uncertainty/RMSE*
Brightness temperature sensitivity(1 s integration time)	*Water-vapour channels ≤ 0.2 K* *Oxygen channels ≤ 0.3 K*
Brightness temperature measurement error	*≤1K RMS*
Relative-humidity-profile RMSE	*≤15%RH*
Temperature-profile RMSE	*H ≤ 2 km: ≤ 1 K* *H > 2 km: ≤ 1.8 K*

Note: The values listed in this table……uses seven channels per band.”

Finally, we also fully agree with the reviewer’s suggestion that a dedicated comparison between stationary land-based deployment and buoy-based deployment would be very valuable. Such an experiment would help separate platform-motion-related effects from the intrinsic performance of the microwave radiometer and would support a more complete characterization of platform-dependent uncertainty. However, this type of dedicated comparison experiment was not included in the present study. Therefore, in the Discussion section of the revised manuscript (Page 23, Lines 575–582), we now explicitly acknowledge this as a limitation of the current work and indicate that such inter-platform comparison experiments will be an important direction for future study. The revised text is as follows: "*A further limitation of the present study is that a dedicated comparison experiment between stationary land-based deployment and buoy-based deployment was not included. Therefore, the current work should not be interpreted as a complete platform-dependent uncertainty characterization of the instrument. Future work will focus on two directions: expanding offshore observational coverage and developing multi-source data-fusion strategies. First, by increasing the number of buoy deployments, a broader marine observation network may be established for more extensive monitoring of atmospheric parameters. Second, buoy observations may be combined with satellite microwave and GNSS remote sensing data to alleviate data scarcity and provide richer information for marine atmospheric analysis and forecasting. In addition, transfer-learning methods may help improve the generalizability of the retrieval framework.*"

Based on these revisions, we have substantially narrowed and clarified the relevant wording in the manuscript. The revised version no longer overstates the degree of instrument-specific characterization of the present QFW-6000 system. Instead, the focus is now clearly placed on demonstrating the feasibility and methodological characteristics of the buoy-based retrieval framework under the tested observational conditions, rather than claiming that a complete metrological characterization of the instrument under different deployment platforms has already been achieved.

Revisions made

1. The uncertainty-related description in Section 2.1 was rewritten to clearly distinguish background performance ranges from instrument-specific characterization.
2. Table 1 was added to summarize the main technical specifications of the QFW-6000 used in this study.
3. The manuscript now explicitly states that the system used here and that in Zhang et al. (2025) belong to the same instrument family but have different channel

configurations.

4. A limitation statement was added in the Discussion to clarify that no dedicated fixed-site versus buoy-based deployment comparison was included in the present study.
5. Related wording throughout the manuscript was tightened to avoid overstating the level of instrument characterization.

Comment L310

L310: What are the bounds? What are the values of delta1 und delta2. If these values are small, this prevents strong inversions. Again, a case study would be helpful to clarify.

Response:

We sincerely thank the reviewer for this important and constructive comment. We fully agree that the original manuscript did not provide sufficient quantitative information on the boundary conditions and the values of δ_1 and δ_2. As the reviewer rightly pointed out, without such information it is difficult for the reader to assess whether the imposed constraints might overly restrict strong inversions, local gradient-transition structures, or sharp layer-to-layer variations. We also agree that a representative case study is helpful, because the practical role of the constraints cannot be fully appreciated from formulas alone.

To address this issue, we made two major additions in the revised manuscript.

First, in the methodological description (Page 14, Lines 325–333), we now explicitly clarify both the mathematical form and the numerical values of the relevant constraints. In the present retrieval framework, the adjacent-layer constraints are expressed as

$$\mid T_{i+1} - T_i \mid \leq \delta_1$$

and

$$\mid H_{i+1} - H_i \mid \leq \delta_2$$

with $\delta_1 = 8\,\mathrm{K}$for temperature and $\delta_2 = 60\%$for relative humidity. The purpose of these constraints is to suppress unrealistically sharp layer-to-layer jumps while preserving basic physical continuity in the retrieved profiles.

At the same time, we further clarify that the lower and upper bounds of the retrieval variables are not fixed universal constants. Instead, they are generated level by level from the month-dependent small-sample prior database. In other words, the inversion is simultaneously controlled by two categories of constraints: (1) level-dependent admissible intervals derived from the prior database, and (2) adjacent-layer continuity constraints. The former restrict the physically plausible search range at each altitude level, whereas the latter suppress unrealistic vertical discontinuities between neighbouring layers.

We agree with the reviewer's concern that if $\delta_1$and $\delta_2$were chosen too small, they could indeed over-suppress local strong inversions or sharp vertical transitions. For this reason, we now clarify more explicitly in the revised manuscript that the present "lapse-rate constraint" is not a strict monotonic decrease constraint with height. Rather, it limits only the absolute magnitude of adjacent-layer differences. Therefore, local inversions or local gradient-transition structures are not excluded a priori, although overly sharp and physically implausible jumps may still be moderated by the imposed limits. We believe this clarification is important to avoid the misunderstanding that the method enforces a strictly monotonic profile shape.

The amended text is as follows: "*In this study, the adjacent-layer limits are set to* $\delta_1 = 8\,K$ *for temperature and* $\delta_2 = 60\%$ *for relative humidity. These constraints are introduced to suppress unrealistically sharp layer-to-layer jumps while maintaining physically reasonable profile continuity. The lower and upper bounds of temperature and relative humidity at each retrieval height are not fixed universal constants; instead, they are generated from the month-dependent small-sample prior database and used as level-dependent admissible intervals for the optimization variables.*

It should be emphasized that the present so-called "lapse-rate constraint" does not enforce a monotonic decrease of temperature with height. Instead, it limits only the absolute magnitude of adjacent-layer changes. Therefore, local inversions are not excluded a priori, and the constraint does not by itself guarantee a fully smooth mean retrieved profile."

Second, following the reviewer's suggestion, we added a representative temperature-profile case study in the Results section (Fig. 10; Page 21–22, Lines 527–534). This example was included specifically to illustrate the practical effect of the imposed bounds and adjacent-layer continuity constraints. The enlarged lower-tropospheric view shows that local gradient-transition features can still be retained under the current constraint settings, indicating that the imposed constraints do not a priori exclude inversion-like or non-monotonic structures. This supports our revised interpretation that the present constraints act primarily as continuity controls rather than as a strict monotonic lapse-rate assumption.

The specific amendments are as follows: "*To further clarify the effect of the imposed bounds and adjacent-layer continuity constraints, a representative case study is presented in Fig. 10.*"

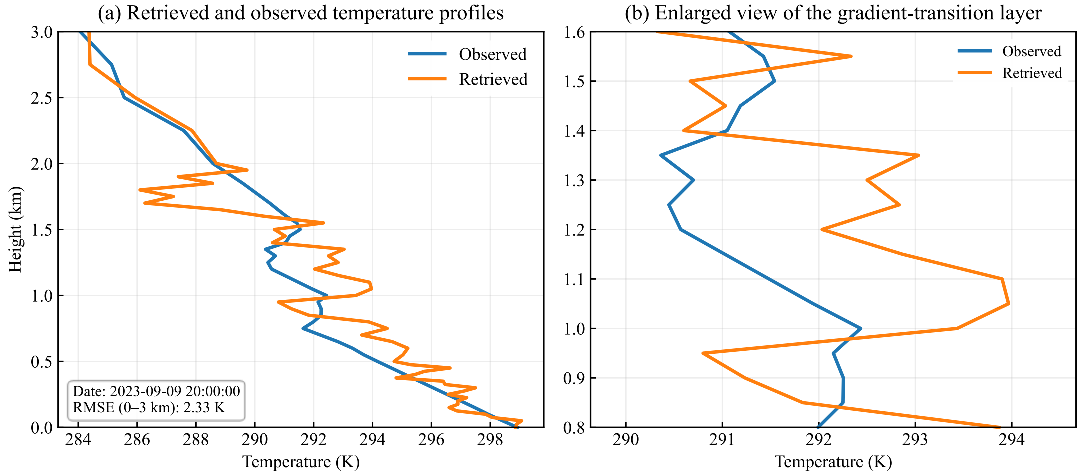

(a) Retrieved and observed temperature profiles

(b) Enlarged view of the gradient-transition layer

Date: 2023-09-09 20:00:00
RMSE (0–3 km): 2.33 K

Through these two additions, we believe that the transparency of the parameter settings and the physical meaning of the imposed constraints have both been improved.

Revisions made

1. The specific values of δ_1 and δ_2 were added in the methodological description.
2. We clarified that the bounds are generated from a month-dependent small-sample prior database and enter the inversion as level-dependent lower and upper limits.
3. We explicitly clarified that the imposed constraints are adjacent-layer continuity constraints rather than strict monotonic decrease constraints.
4. A representative case study (Fig. 10) was added to illustrate the practical influence of the imposed bounds and continuity constraints on the retrieved profile structure.

Comment - Results in Figure 8: The results are based on a systematic error correction using 38 radiosondes. These sondes are representative for the Jiaozhou Gulf. There is no variation or seasonality in the dataset. Additionally, all radiosondes were launched within one month. How do the conditions differ? I can't see how independent the test data set is. I mean, how independent is the data if the conditions could be almost identical? Please comment on that and describe the variability in the radiosonde dataset.

I strongly doubt that the accuracy shown here in Figure 8 can be achieved on the open ocean. To justify its use on the open ocean and to fit the motivation, the inaccuracy analysis based on 38 radiosondes, which were also collected at a land

station, seems too weak to me.

Response:

We sincerely thank the reviewer for this important and highly relevant comment. We fully agree that the original manuscript did not provide a sufficiently clear description of the variability, representativeness, and limitations of the 38 matched radiosonde cases used for the systematic error correction analysis. We also agree with the reviewer's assessment that the previous wording could leave the impression that the accuracy shown in the original Figure 8 could be directly generalized to broader offshore or open-ocean conditions. Upon reconsideration, we recognize that this interpretation was indeed too strong and insufficiently cautious.

In response, we made two major revisions.

First, in the Results section of the revised manuscript (Page 20–21, Lines 509–526), we now define much more clearly the representativeness boundary of this dataset and the appropriate scope of interpretation of the corresponding retrieval results. We explicitly state that these 38 matched radiosonde profiles were collected during a relatively short campaign period and therefore mainly represent the nearshore atmospheric conditions encountered during the Jiaozhou Bay experiment. They do not represent the full range of possible marine atmospheric states and should certainly not be equated with open-ocean conditions. We now clearly state that the present dataset does not provide seasonal coverage and does not constitute a climatologically representative sampling of marine boundary-layer variability.

At the same time, we also expanded the description of the variability within the dataset itself. Although all matched radiosonde cases were obtained within one month, they are not strictly identical repetitions under the same conditions. In the revised manuscript, we now clarify that the dataset includes both daytime and nighttime radiosonde launches and that there is still case-to-case variability in near-surface thermodynamic conditions, lower-tropospheric temperature–humidity stratification, and local gradient-transition structures. In other words, although the dataset is clearly limited in temporal span and environmental representativeness, it should not be viewed as a collection of perfectly repeated cases under nearly identical atmospheric conditions.

Second, in response to the reviewer's concern regarding applicability to open-ocean environments, we substantially narrowed the scope of the corresponding conclusions. Specifically, in the revised Results section (Page 20–21, Lines 509–526), we now explicitly state that the result originally shown in Figure 8, and the corresponding revised figure, should be interpreted primarily as evidence of retrieval feasibility and systematic-error characterization under the tested nearshore observational conditions, rather than as a generalized validation at the same level of accuracy for broader offshore or open-ocean environments. In other words, the present analysis, based on 38 matched radiosonde profiles and a correction framework tied to nearshore and land-station-based collocation, is not sufficient to support strong general claims regarding accuracy in open-ocean conditions.

The amended text is as follows: "*Despite these physical limitations, the RMSE values still indicate that the proposed method can provide operationally useful humidity information under the present nearshore observational conditions. Consistent with previous studies using ground-based microwave radiometers, retrieval RMSE values on the order of 1–2 K for temperature and 10–30% for relative humidity are generally regarded as operationally useful (Massaro et al., 2015; Yan et al., 2020; Cimini et al., 2011).*

……

Accordingly, the results in Fig. 9 should be interpreted primarily as evidence of retrieval feasibility and systematic-error characterization under the tested nearshore buoy-based conditions, rather than as a generalized validation across all marine environments."

In addition, we further revised the Discussion section (Page 22–23, Lines 558–574) to state explicitly that extending the present framework to broader nearshore and open-ocean applications will require additional validation using marine observational datasets with much greater temporal variability and environmental diversity.

The amended text is as follows: "*The sea-trial results obtained in Jiaozhou Bay demonstrate the effectiveness and feasibility of the proposed method under the present campaign conditions.*

……

Extension of the present framework to broader offshore and open-ocean applications will require further validation using more diverse marine datasets with larger temporal and environmental variability."

We believe that this revision more appropriately defines the valid scope of the current conclusions and avoids overextending what can be supported by the present dataset.

Revisions made

1. We clarified the temporal coverage and representativeness boundary of the 38

matched radiosonde profiles in the Results section.

2. We added an explanation that the dataset includes both daytime and nighttime launches and still exhibits case-to-case variability in near-surface thermodynamic conditions and lower-tropospheric structures.
3. We explicitly state that the corresponding result should be interpreted as representing the present nearshore experimental conditions rather than a general validation across all marine environments.
4. We narrowed the wording concerning open-ocean applicability and recast it as a question requiring future validation rather than as an already established conclusion.
5. We added an explicit limitation statement in the Discussion emphasizing that broader marine applicability requires additional validation with more diverse datasets.

Response to RC2

Comment on In the response, the authors state that the reason for the oscillations in the average retrieved temperture profile are the product of low weighting function sensitivity leading to 'stratification structures ... be[ing] retained'. They then state this behaviour is typical of microwave radiometry. This is not typical of microwave radiometer retrievals, and by my own experience, the opposite is true about low weighting function sensitivity- that this leads to smoother retrieved profiles as information is spread across large altitude differences. I suspect, that the oscillations in the mean temperature is due to insufficient regularisation of the algorithm. I also find these oscillations in the mean temperature profile hard to understand with the statement that the algorithm has a lapse rate constraint in it, which presumably should prevent temperature increases from a lower level to a higher level.

Response：

We sincerely thank the reviewer for this technically insightful and physically important comment. We fully agree that the original wording, which described the oscillatory behaviour in the mean retrieved temperature profile as a "typical feature" of microwave radiometer retrievals, was not sufficiently rigorous and could lead to misunderstanding regarding the physical origin of this behaviour.

To address this concern, we revised both the interpretation of the oscillations and the description of the imposed constraints.

First, in the Results section (Page 17, Lines 415–425), we removed the wording suggesting that the oscillatory behaviour is a "typical feature" and replaced it with a more cautious interpretation. In the revised manuscript, we now explain that the observed oscillations are more appropriately understood as the combined result of reduced information content in altitude ranges where the microwave weighting functions are less sensitive, the structural characteristics of the discrete layer-by-layer optimization framework, and the current regularization setting. We also agree with the reviewer that insufficient regularization may be a contributing factor to the observed oscillatory behaviour. Through this revision, we avoid attributing the phenomenon to an intrinsic or universal characteristic of microwave radiometer retrievals.

The specific amendments are as follows: "*The radiosonde profiles shown here represent ensemble-averaged means over 38 matched cases, which suppress small-scale vertical variability and therefore appear smoother.*

……

In particular, this behaviour may also indicate that the current regularization strength is not fully sufficient to suppress small-scale layer-to-layer oscillations under conditions of limited observational constraint. Further refinement of the regularization strategy may help improve the smoothness and physical consistency of the retrieved profiles."

Second, in the methodological description (Page 14, Lines 325–333), we further clarified the actual physical meaning of the so-called "lapse-rate constraint." In the revised manuscript, we explicitly state that this constraint is implemented as an adjacent-layer continuity constraint rather than as a strict monotonic decrease constraint. Since it limits only the magnitude of adjacent-layer differences and not their sign, it does not exclude local inversions or gradient-transition structures a priori. This clarification directly addresses the reviewer's concern as to why oscillatory or locally non-monotonic structures can still appear despite the presence of the so-called "lapse-rate constraint."

The specific amendments are as follows: "*In this study, the adjacent-layer limits are set to* $\delta_1 = 8K$ *for temperature and* $\delta_2 = 60\%$ *for relative humidity. These constraints are introduced to suppress unrealistically sharp layer-to-layer jumps while maintaining physically reasonable profile continuity. The lower and upper bounds of temperature and relative humidity at each retrieval height are not fixed universal constants; instead, they are generated from the month-dependent small-sample prior database and used as* *level-dependent admissible intervals for the optimization variables.*

It should be emphasized that the present so-called "lapse-rate constraint" does not enforce a monotonic decrease of temperature with height. Instead, it limits only the absolute magnitude of adjacent-layer changes. Therefore, local inversions are not excluded a priori, and the constraint does not by itself guarantee a fully smooth mean retrieved profile."

In addition, following the reviewer's suggestion, we added a representative case study (Fig. 10) in the Results section (Page 21–22, Lines 527–534). This example is intended to more directly illustrate the practical role of the imposed bounds and continuity constraints. The enlarged lower-tropospheric view shows that local gradient-transition features can still be retained in the retrieved profile, indicating that the imposed constraints do not enforce a strict monotonic decrease of temperature with height. This supports our revised interpretation that the current constraints act mainly as continuity controls rather than as overly restrictive smoothing operators.

The specific amendments are as follows: "*To further clarify the effect of the imposed bounds and adjacent-layer continuity constraints, a representative case study is presented in Fig. 10.*"

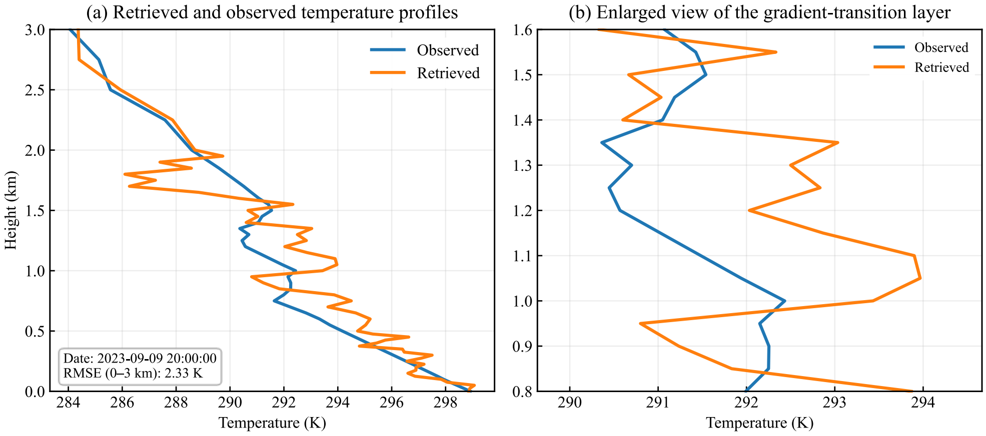

Furthermore, in the Discussion section of the revised manuscript (Page 23, Lines 578–581,Lines 587-593), we have also rewritten the relevant discussion to more explicitly acknowledge that the current retrieval framework still has room for improvement. In particular, we now state that the existing atmospheric prior database and multi-objective optimization strategy provide preliminary and effective constraints, but that more comprehensive physical constraints and more systematic optimization of the evolutionary parameters may further reduce retrieval uncertainty and improve both retrieval accuracy and computational efficiency.

Taken together, these revisions provide a more physically consistent and logically coherent explanation of the oscillatory behaviour and fully incorporate the reviewer’s important suggestion regarding the possible role of insufficient regularization. The specific revised text is as follows:*These mid-tropospheric fluctuations are more appropriately interpreted as manifestations of limited information content in the weighting-function trough region, combined with the present regularization setting, rather than as a typical or unavoidable characteristic of microwave radiometer retrievals in general. Nevertheless, the proposed method still demonstrates significant application potential under the tested conditions. In terms of algorithm optimization, the atmospheric prior experience database and multi-objective optimization strategy (NSGA-II) provide preliminary and effective constraints, but there is still room for further improvement. The current retrieval framework includes boundary constraints and adjacent-layer continuity constraints; Future research may explore the introduction of more comprehensive physical constraints to further reduce retrieval uncertainty and improve accuracy. In addition, the settings of evolutionary parameters in NSGA-II have a critical impact on convergence speed and global optimization capability. Subsequent work may employ more systematic simulation-based optimization strategies to fine-tune these parameters, with the aim of achieving higher retrieval accuracy and computational efficiency."*

Taken together, these revisions provide a more physically consistent and logically coherent explanation of the oscillatory behaviour and fully incorporate the reviewer’s important suggestion regarding the possible role of insufficient regularization.

Revisions made

1. We removed the wording that described the oscillatory behaviour as a “typical feature” and replaced it with a more cautious physical interpretation in the Results section.
2. We explicitly clarified in the Method section that the so-called “lapse-rate constraint” is in fact an adjacent-layer continuity constraint rather than a strict monotonic decrease constraint.
3. We revised the Discussion to provide a unified and more careful interpretation of the remaining fluctuations in the middle troposphere.
4. We added a representative case study (Fig. 10) to illustrate the practical effect of the imposed bounds and continuity constraints on the retrieved profile structure.

Comment

Question about improvement with information from attitude handling

I asked about the difference that the attitude correction makes compared to handling the data with no attitude correction. I agree with the authors that the pointing angle of the radiometer is important in the retrieval step, and so a quantification of this improvement would be a very useful output of the paper. As the paper stands, it seems there is a part missing about why their algorithm (which is able to take in information about the platform) should be used as opposed to other algorithms without this. I would encourage the authors to make a comparison at least to the radiometer's own retrieval (presumbly they have this as they state that the LWP was found from this).

Response:

We sincerely thank the reviewer for this important and practically meaningful suggestion. We fully agree that, since one of the key features of the proposed framework is its explicit incorporation of platform attitude information into the retrieval process, it is necessary to quantitatively assess the difference between retrievals with and without attitude correction and to compare the proposed framework with the built-in radiometer retrieval product in order to better justify its practical value.

To address this comment, we first added a quantitative comparison between retrievals with and without attitude correction in the revised manuscript (Page 17–19, Lines 443–476). In this comparison, the same brightness temperature observations are used in all cases, and the only difference is whether platform attitude information is incorporated into the forward radiative transfer model. To evaluate this effect more systematically, we not only present the overall comparison but also add grouped statistics for different attitude-angle ranges and a representative case with relatively large attitude variation.

These results show that, for the present matched dataset, the overall average improvement associated with attitude correction is relatively modest. For example, in the representative case, the temperature RMSE decreases from 3.69 K to 2.99 K, while the relative humidity RMSE decreases from 18.95% to 18.57%. The grouped statistics also show that the positive effect of attitude correction tends to become more evident in the subset with larger attitude angles, whereas the average effect across the entire dataset remains limited.

We then added a physical explanation of this apparently modest average improvement. In the revised manuscript, we argue that the platform attitude angles in the present nearshore dataset are generally small, typically around 1–3°. Under such weak-tilt conditions, the geometric perturbation to the observation path is relatively small, and its effect can be partially masked by other sources of retrieval uncertainty. To further examine the physical necessity of attitude correction, we supplemented the manuscript with a theoretical sensitivity analysis. The results show that changes in viewing angle can introduce systematic brightness-temperature deviations across multiple K-band and V-band channels. This means that, even though the average improvement is limited for the present dataset, attitude correction remains physically necessary from the perspective of radiative transfer and is expected to play a more significant role under stronger platform-motion conditions and larger viewing-angle deviations.

The specific amendments are as follows: "*The retrieval results shown in Fig. 7 already incorporate attitude correction in the forward modelling process. To isolate the effect of attitude correction, the following analysis is conducted by additionally comparing retrievals with and without applying the attitude correction, before the system error correction is introduced.*

……

The improvement in relative humidity is less consistent than that in temperature, which may be attributed to the higher sensitivity of humidity retrieval to measurement noise and vertical smoothing, as well as the comparatively weaker viewing-angle sensitivity of the humidity-sensitive channels.

Table 2. *Grouped mean changes in retrieval RMSE under different attitude-angle ranges.*

Attitude angle range	*Mean ΔRMSE(After-Before,T,K)*	*Mean ΔRMSE(After-Before,RH,%)*
≤1.5°	*-0.14*	*0.18*
1.5–2.5°	*-0.04*	*0.06*
>2.5°	*0.13*	*0.07*

For a representative case with relatively large attitude variation, the quantitative comparison before and after attitude correction is summarized in Table 3. The temperature RMSE decreases from 3.69 K to 2.99 K, while the relative humidity RMSE decreases from 18.95% to 18.57%. These results indicate that incorporating attitude information improves the physical consistency of the forward model and becomes increasingly important under conditions with stronger platform motion.

Table 3. *Retrieval performance before and after attitude correction for a representative case.*

Case time	*Angle (°)*	*Temperature RMSE (K)*			*Relative humidity RMSE (%)*		
		Before	*After*	*Reduction*	*Before*	*After*	*Reduction*
2023-09-22 20:00	*2.91*	*3.69*	*2.99*	*0.70*	*18.95*	*18.57*	*0.38*

……

Figure 8 was generated from a theoretical sensitivity experiment in which a standard atmospheric profile from ITU-R P.835 was combined with ITU-R P.676 gaseous absorption calculations to simulate channel brightness temperatures at different viewing angles; ΔTB was defined relative to the 0° nadir-view case.

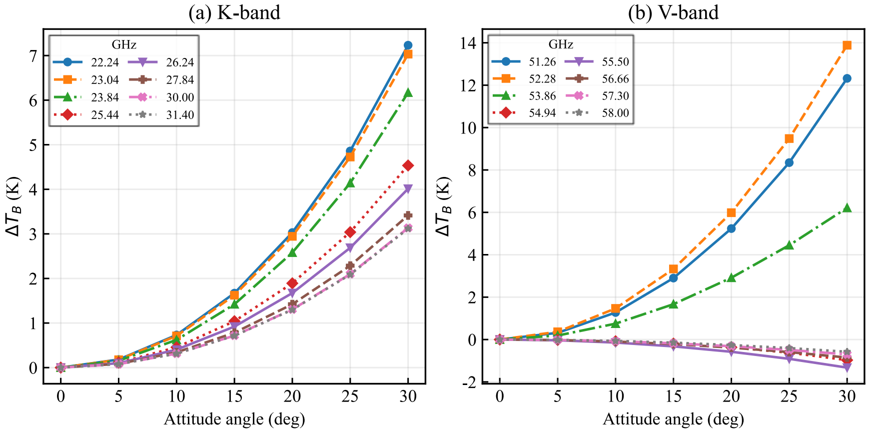
(a) K-band

(b) V-band

”

Second, following the reviewer’s recommendation, we added a direct quantitative comparison with the built-in radiometer retrieval algorithm (Page 22, Lines 535–547; Table 5). The results show that, for temperature, the built-in algorithm yields an RMSE of 4.13 K, while the proposed method achieves 4.11 K before correction and 2.08 K after correction. For relative humidity, the built-in algorithm yields an RMSE of 29.50%, whereas the proposed method achieves 24.09% before correction and 20.95% after correction. These results indicate that the proposed framework maintains comparable baseline performance and achieves improved retrieval accuracy after correction for the present matched dataset.

The specific amendments are as follows: “*To further evaluate the performance of the proposed method, a quantitative comparison with the built-in retrieval algorithm of the microwave radiometer is conducted over the 0–10 km altitude range.*

All methods are evaluated on the same matched samples and interpolated onto a common vertical grid to ensure consistency. The comparison results are summarized in Table 5.

Table 5. *Comparison of retrieval performance between the proposed method and the built-in radiometer algorithm (0–10 km).*

Method	*Temperature RMSE (K)*	*Relative Humidity RMSE (%)*
Built-in Radiometer Algorithm	*4.13*	*29.50*
Proposed Method (Before Correction)	*4.11*	*24.09*
Proposed Method (After Correction)	*2.08*	*20.95*

……

Overall, the results indicate that the proposed method maintains a comparable baseline performance to the built-in algorithm and achieves improved retrieval accuracy after correction for the present matched dataset."

More importantly, unlike the built-in algorithm, the present framework explicitly incorporates platform attitude information into the forward modelling process and therefore provides a more physically consistent treatment of the observation geometry. This feature is particularly relevant for buoy-based deployment scenarios in which platform motion cannot be neglected. Through these revisions, we believe that the manuscript now more clearly demonstrates both the necessity of attitude-aware processing and the practical added value of the proposed method.

Revisions made

1. We added a direct comparison between retrievals with and without attitude correction to quantify the effect of incorporating platform attitude information.
2. We added grouped statistics and a representative case to show how the attitude effect depends on the magnitude of platform motion.
3. We added a direct comparison with the built-in radiometer retrieval product (Table 5).
4. We supplemented the text with a physical explanation of why the average improvement remains limited under the present matched-sample conditions.

Comment: Unanswered Comments in General comments section

I made the comment in the general comments: "When the zenith angle of the radiometer is under constant flux as it is on a sea surface, the integration time of a single observation is very relevant, yet it is not discussed in the paper."

I think this would be something relevant to either the technical details or the discussion part of the paper. the attitude sensor gathers data every second, but the integration time of the radiometer should be much longer. So how are the many data point from the attitude sensor passed to the algorithm? Is it kept as it is and the algorithm calculates a sort of effective zenith angle? Or is there some preprocessing to calculate this manually?

Response:

We thank the reviewer for this important technical comment. We fully agree that the temporal consistency between the high-frequency attitude measurements and the radiometer integration time must be explicitly clarified in the methodological description.

In the revised manuscript (Page 12–13, Lines 293–298), we added a detailed explanation of this processing step. Specifically, the attitude sensor samples at a higher frequency (approximately 1 Hz or higher), whereas each microwave radiometer observation represents a time-integrated brightness temperature over a finite observation window. To ensure temporal consistency, all attitude samples within each radiometer integration window are first synchronized and averaged to obtain an effective attitude state. Based on this averaged attitude, an effective zenith angle is then calculated and used in the forward radiative transfer model.

The specific amendments are as follows: "*Because the attitude sensor samples at a higher frequency than the microwave radiometer, the raw attitude measurements are first synchronized to the radiometer integration period. All attitude samples within one radiometer observation window are averaged to derive an effective attitude state. Based on this averaged attitude, an effective zenith angle is calculated and used in the forward radiative transfer model. In this way, the observation geometry remains temporally consistent with the integrated brightness temperature measurement. This treatment ensures consistency between the time-integrated radiometric observation and the observation geometry used in the retrieval.* "

This treatment is adopted because the radiometric observation itself is a time-integrated quantity, and therefore the corresponding observation geometry should also be represented in a time-averaged sense. In addition, this approach avoids introducing high-frequency attitude fluctuations directly into the retrieval process, since such fluctuations are not resolved by the radiometer measurement itself.

By adding this clarification, we believe that the methodological description has become both clearer and more reproducible.

Revisions made

1. We added an explicit description of the temporal matching between attitude data and radiometer integration periods in the Method section.
2. We clarified that the effective attitude state is obtained by averaging the attitude samples within each radiometer integration window.

3. We clarified how the effective zenith angle is calculated and how it is used in the forward radiative transfer model.

Additional correction:

During the review process, we discovered that one reference entry in the previous version of the manuscript had been inadvertently truncated. This issue has been corrected on page 28 of the revised manuscript, lines 739–741, and the entire reference list has been carefully checked to ensure its completeness and consistency.