Response to Reviewer #2

Major Comments:

1. The theoretical interpretation of adaptive observation weights requires further clarification.

The central innovation of this study is the introduction of adaptive observation weights into the TCKF1D-Var cost function. However, the current manuscript does not sufficiently explain the theoretical meaning of these optimized weights. According to Eqs. (3)–(5), the observation weights are introduced as adjustable parameters during the minimization process. However, it remains unclear whether these weights represent the relative information contributions of different observations or the mathematical scaling factors introduced to optimize retrieval performance.

The manuscript frequently interprets large weights as indicating higher observation information content. However, the relationship between optimized weights and formal information content is not rigorously demonstrated. For example, the manuscript states that certain GMWR channels with large weights provide dominant observational constraints and that higher-altitude MRL observations receive larger confidence. However, because the weights are determined jointly by observation residuals, background uncertainties, and the retrieval framework itself, these interpretations may not necessarily indicate the intrinsic information content of the observations.

The authors should clarify the physical interpretation of the adaptive weights, whether the diagnosed weights are comparable among different instruments and observation types, and what assumptions are required when interpreting weights as information contributions.

> **Re:**
>
> We appreciate the reviewer poniting out this insightful comment. We agree that the theoretical meaning of the adaptive observation weights was not sufficiently clarified in the original manuscript. In particular, our previous discussion may have overstated the relationship between the optimized weights and the information content of individual observations. We have therefore clarified that the optimized observation weights should primarily be interpreted as adaptive mathematical weighting factors that regulate the relative influence of different observations within the retrieval framework, rather than as direct measures of their intrinsic or formal information content (**Lines 175–180**).
>
> According to Eqs. (3)–(5), the observation weights are introduced as adjustable parameters during the minimization process. Their mathematical role is to modify the relative contribution of the observation-related terms to the retrieval cost function, thereby allowing the retrieval system to adaptively balance observational and background information. The optimized weights are determined by the joint interaction among the observation residuals, background uncertainties, retrieval constraints, and the structure of the retrieval framework. Therefore, a larger optimized weight does not necessarily indicate that the corresponding observation intrinsically contains more information. Instead, it indicates that, under the specific atmospheric conditions and within the specific retrieval configuration considered here, greater relative influence is assigned to that observation in constraining the retrieved atmospheric state.

Regarding the relationship between the optimized weights and observational information content, we agree that the current study does not provide sufficient evidence to establish a quantitative or one-to-one correspondence between the two. Nevertheless, the optimized weights can provide useful diagnostic information about the relative observational influence of different channels, range bins, and observing systems. For example, systematic variations in the weights may reflect how the retrieval system adjusts the contribution of different observations in response to changes in observation–background consistency, observational uncertainty, and atmospheric conditions. Such behavior may contain characteristics associated with the relative information contributions of different observations, but these characteristics should be regarded as indirect rather than formal measures of information content. Establishing the quantitative relationship between optimized weights and formal information content would require additional diagnostics, such as controlled sensitivity experiments or information-content-based analyses, which are beyond the scope of the present study. The corresponding revisions have been incorporated into the revised manuscript at several relevant locations (**Lines 72, 84–86, 300–303, 323–327, 437–440, 660–663, 697, and 716–727**).

We also agree that the comparability of the optimized weights among different instruments and observation types needs to be carefully qualified. The absolute values of the weights should not be interpreted as directly comparable measures of observational information across GMWR channels, MRL range bins, or different observing systems, because their optimized values depend on the error characteristics, observation operators, normalization of the cost function, and the role of each observation type in the specific retrieval framework. Comparisons of the weights are therefore most meaningful when the observations are considered within the same retrieval formulation and under consistent assumptions regarding observational and background uncertainties. In the revised manuscript (**Lines 180–187**), we have consequently emphasized comparisons of the relative observational influence of observations within the same retrieval configuration, rather than interpreting the absolute magnitude of a weight as a universal measure of information content.

2. **The validation strategy should better separate retrieval improvement from possible effects of the background field.**

The manuscript evaluates retrieval performance against radiosonde observations and compares adaptive weighting with static weighting. This approach is reasonable. However, because ERA5 profiles are used as prior information, the retrieved profiles may partly inherit characteristics from the background field. The manuscript should provide more discussion regarding whether adaptive weighting mainly adjusts observation influence or compensates for deficiencies in ERA5, and whether the improvements are consistent under different background uncertainties.

Re:

Thank you for this insightful comment. We agree that the use of ERA5 as the background state introduces an additional source of information into the retrieval and that the potential influence of the background field should be discussed. In the revised manuscript (**Lines 147–156**), we have clarified the role of ERA5 and the mechanism through which adaptive observation weighting affects the balance between observational and background information.

In the present TCKF1D-Var framework, ERA5 provides the a priori atmospheric state and contributes to the background term of the cost function. Consequently, the retrieved profiles can retain characteristics of the ERA5 background, particularly when the observational constraints are weak or when the assumed background uncertainty is relatively small. The adaptive observation weighting is not specifically designed to compensate for deficiencies or biases in ERA5. Rather, its primary role is to dynamically regulate the relative influence of the observations and the background information during the retrieval process. When the observations provide stronger constraints, the adaptive weighting allows them to exert greater relative influence, thereby reducing the dependence of the retrieval on the background state. Conversely, when the observational constraints are weaker, the background information retains a larger role in determining the retrieval solution. Therefore, the improvements observed in the adaptive-weighting experiments should be interpreted primarily as a consequence of a more flexible balance between observations and background information, rather than as direct evidence that the method corrects deficiencies in ERA5. The comparison with radiosonde observations indicates that the adaptive weighting provides more accurate retrievals than the static-weighting configuration under the cases investigated in this study. However, we acknowledge that this comparison alone cannot fully separate the contribution of improved observation weighting from the influence of the ERA5 background.

Regarding the dependence on background uncertainties, we agree that a systematic sensitivity analysis using different background-error specifications would provide a stronger assessment of the robustness of the proposed method. Such an analysis is beyond the scope of the present study, which uses a single ERA5 background dataset and a fixed background-error specification. We have therefore explicitly acknowledged this limitation in the revised manuscript (**Lines 742–745**) and identified sensitivity experiments with different background uncertainties as an important direction for future work.

3. The physical interpretation of precipitation-dependent performance requires further discussion.

The manuscript shows that the advantage of adaptive weighting becomes weaker under precipitation exceeding 30 mm/hr, especially for GMWR-only retrievals. This is an interesting result; however, the explanation that increased observational uncertainty reduces the effectiveness of adaptive weighting remains speculative. The authors should clarify whether precipitation contamination effects were explicitly considered in the observation processing, and whether the signal attenuation and cloud effects influence the results.

Re:

We agree that the explanation of the reduced advantage of adaptive weighting at hourly precipitation rates exceeding 30 mm, particularly for the GMWR-only retrievals, should be stated more cautiously. We have therefore revised the relevant discussion (**Lines 337–345**) and clarified how precipitation-related effects were controlled in the observation processing and case analysis.

First, precipitation contamination of the GMWR observations was explicitly controlled using the collocated minute-level precipitation observations. During the retrieval preprocessing, GMWR observations corresponding to periods with precipitation were identified and excluded from the retrieval. Therefore, the GMWR observations used in the analyzed retrievals do not directly correspond to periods of active precipitation. Second, for the analysis of the cases with intense precipitation, we selected the retrieval profiles temporally closest to the occurrence of the intense precipitation events. This design allows us to examine the retrieval behavior immediately surrounding the strong-precipitation events while avoiding the direct inclusion of rain-contaminated GMWR measurements. Taken together, these procedures minimize the direct influence of precipitation contamination on the GMWR observations to the greatest extent possible within the present experimental framework. Therefore, the weaker advantage of adaptive weighting at precipitation rates exceeding 30 mm should not simply be interpreted as a consequence of directly contaminated GMWR observations.

Nevertheless, we agree that precipitation-related effects cannot be considered completely negligible. Even when GMWR observations during precipitation are excluded, intense precipitation events are generally accompanied by rapid changes in cloud liquid water, atmospheric moisture, and hydrometeor distributions, which may indirectly affect microwave radiometric observations and the consistency between the observations and the retrieval framework. In addition, MRL observations may be affected by precipitation-related signal attenuation and hydrometeor scattering. These effects were not independently quantified in the present study. Consequently, our revised manuscript no longer attributes the reduced effectiveness of adaptive weighting solely to increased observational uncertainty. Instead, we describe it as a possible combined effect of rapidly evolving atmospheric conditions and residual precipitation-related influences on the observational constraints.

Minor Comments:

1. Clarify the novelty compared with previous TCKF1D-Var studies.

The manuscript extends previous studies by Zhang et al. (2026a,b). However, the distinction between the previous TCKF1D-Var framework and the current adaptive weighting version should be highlighted more clearly in the Introduction. A concise comparison table showing the differences between previous method and current improvement would help readers understand the methodological contribution.

Re:

We agree that the methodological distinction between the previously developed TCKF1D-Var frameworks and the adaptive observation weighting version proposed in this study was not sufficiently emphasized in the original Introduction. The previous GMD study established the TCKF1D-Var framework for GMWR observations by introducing the virtual potential temperature constraint, a ratio-based cost function, and a cloud-microphysics closure to improve the physical consistency and stability of thermodynamic and hydrometeor retrievals. The subsequent AMT study extended this framework to Mie–Raman lidar observations by introducing a physics-informed Raman lidar observation operator and extending the framework to high-resolution water-vapor profiling.

Following the reviewer's suggestion, we have therefore added a concise methodological comparison in **Appendix A (Table A1, referenced in Line 81)**, which summarizes the differences among the original GMWR-based TCKF1D-Var framework, the subsequently extended MRL-based framework, and the adaptive observation weighting TCKF1D-Var proposed in this study. The table compares their observation types, retrieval variables, physical constraints, observation operators, weighting strategies, treatment of observational uncertainties, and major methodological advances.

2. Terminology of "dynamic weight" and "adaptive weight".

The manuscript uses both "adaptive observation weighting" and "dynamic weight" to describe the proposed method. Please consider using one consistent terminology throughout the manuscript. "Adaptive observation weighting" appears more appropriate because the method adjusts weights during optimization rather than necessarily evolving with time.

Re:

The "dynamic weight" are changed to "adaptive weight" throughout the manuscript.

3. Definition of observation weights should be added explicitly.

The manuscript should provide a clear mathematical definition of the weight coefficient. For example, Is a weight of 1 equivalent to the original observation contribution? Does a weight smaller than 1 indicate reduced confidence? Can weights larger than 1 occur? These details are important for reproducibility.

Re:

We agree that the mathematical definition and interpretation of the adaptive observation weight coefficient should be stated more explicitly to ensure the reproducibility of the proposed method. In the revised manuscript, we have clarified that the weight coefficient is introduced as a multiplicative scaling factor for the corresponding observation term in the retrieval cost function. Specifically, a valule of 1 indicates that the observation retains its original contribution in the unweighted formulation. A value between 0 and 1 reduces the relative contribution of the corresponding observation to the cost function. These clarifications have been incorporated into the revised manuscript in the description of the adaptive observation weighting formulation (**Line 172**).

4. Improve description of MRL observation quality control.

The manuscript uses MRL observations under heavy precipitation conditions. Since Raman lidar performance can be strongly affected by clouds and precipitation attenuation, the manuscript should provide more details about the cloud screening, the missing data treatment, and the signal quality criteria.

Re:

We agree that the quality control of MRL observations is particularly important for retrievals under heavy-precipitation conditions, given the potential impacts of clouds, precipitation attenuation, and signal degradation on Raman lidar measurements. The detailed quality-control procedure for the MRL observations used in this study, including cloud screening, missing-data treatment, and signal-quality criteria, has been described and evaluated in detail in our previous work (Zhang et al., 2026b).

In the present study, we follow the same MRL data-processing and quality-control procedures as those established in Zhang et al. (2026b), including the identification and removal of observations affected by insufficient signal quality and the treatment of missing or unreliable retrieval bins. Therefore, the MRL profiles used in this study have undergone the established quality-control procedure before being incorporated into the TCKF1D-Var retrieval framework. We have added an explicit statement in the revised manuscript referring readers to Zhang et al. (2026b) for the detailed description of the MRL quality-control methodology, while clarifying its application in the present study.

We also acknowledge that, under heavy precipitation, cloud and precipitation-related effects may still influence the availability and quality of MRL observations even after quality control. Therefore, the quality-control procedure should be understood as a means of minimizing the use of unreliable observations rather than completely eliminating all precipitation- and cloud-related effects.

5. **Clarify the retrieval vertical resolution.**

The manuscript discusses different vertical resolutions of GMWR and MRL observations, but the retrieval framework appears to use a common vertical grid. Please clarify the vertical resolution of the final retrieved profiles and how vertical resolution mismatch is handled.

Re:

We deeply agree that the vertical resolution of the final retrieved profiles and the treatment of the different native vertical resolutions of the GMWR and MRL observations should be clarified. In the present study, the vertical grid of the retrieved profiles is consistent with that adopted in Zhang et al. (2026b). Specifically, the retrieval domain extends from the surface to 10.2 km above ground level, with a vertical resolution of 30 m from 0 to 3 km, 150 m from 3 to 6 km, and 300 m from 6 to 10.2 km. Therefore, the final retrieved thermodynamic profiles are reported on this non-uniform vertical grid. The vertical resolution has been clarified in the revised manuscript (**Lines 184–186**).

6. **Some statements in the conclusion imply that adaptive weights reveal physical observation information content.**

Since this interpretation is not directly validated through formal information-content analysis, the wording “information content ” should be moderated in Lines 74, 81, 287, 309, 422, 645, 679.

Re:

Changed.

7. **Line 19, “dynamically estimates observational contributions” can be replaced with “dynamically estimates the relative contribution of individual observations” for better understanding.**

Re:

Replaced.

8. Lines 22–23, “with the largest benefits found for water vapor mass mixing ratio profiles” should be “with the most significant improvements found for water vapor mass mixing ratio profiles”.

Re:

Replaced.

9. Line 52, “a MRL” should be “an MRL”.

Re:

Corrected.

10. Line 114, “The MRLs emits” should be “The MRLs emit”.

Re:

Corrected.

11. Line 197, “sample amount” should be “sample size”.

Re:

Corrected.

12. Lines 203–204, “leads to a improved retrieval performance” should be “leads to an improved retrieval performance”.

Re:

Corrected.

13. Lines 206 and 207, “framewrok” should be “framework”.

Re:

Corrected.

14. Lines 226–227, “prior thermodynamic profiles generated from ERA5 reanalysis” should be “prior thermodynamic profiles generated from the ERA5 reanalysis”.

Re:

Corrected.

15. Table 1, “Case Amount” should be “Number of Cases”.

Re:

Changed.

16. **The title of Figure 11b should be "Hourly Precipitation ∈ [10, 20) mm", please make corrections.**

Re:

Corrected.

17. **Lines 462–463, "The fact that the adaptive observation weighting framework consistently achieves" can be changed to "The consistent lower retrieval errors achieved by the adaptive framework suggest that" for better understanding.**

Re:

Changed.