Response to the Referee's Comments

November 05, 2026

I thank the authors for their detailed responses to my previous comments and for the substantial revisions made to the manuscript. The additional sensitivity analyses, methodological clarifications, and more explicit discussion of the retrieval limitations have improved the manuscript considerably. Most of my previous concerns have been adequately addressed. I have a few remaining points, mainly concerning the interpretation and presentation of the retrieval results, which I believe can be resolved through minor clarifications.

Response: We thank the referee for the constructive assessment of the revised manuscript and for these helpful suggestions. We address each point below.

Comment 1: I appreciate the authors' clarification in Response 1 that ($C_{\mathrm{eff}} = 5$) is a reference scenario rather than an independently constrained value. However, the abstract still reports single mean values of (k), SSA, and LR without explicitly linking them to this assumption. Figure 5 shows substantial sensitivity, with median LR_{355} varying from 137.3 to 84.3 sr and SSA_{355} from 0.646 to 0.867 as (C_{eff}) increases from 1.57 to 5. I suggest explicitly identifying the reported values as conditional on ($C_{\mathrm{eff}} = 5$) in the abstract and including a representative sensitivity range, particularly for LR_{355}. This would make the presentation consistent with the authors' revised interpretation and avoid implying that the reported values are uniquely constrained by the observations.

Response: We agree. We have placed the $C_{\mathrm{eff}} = 5$ reference-scenario condition before the numerical results in the abstract and added the representative sensitivity reported in Fig. 5. The abstract explicitly distinguishes the mean values obtained under the reference scenario from the scenario-dependent medians used to illustrate sensitivity to C_{eff}.

Changes in the manuscript:

Accurate constraints on aerosol complex refractive index (CRI), single scattering albedo (SSA), and lidar ratio (LR) are essential for improving lidar-based aerosol retrievals and reducing radiative-effect uncertainties. However, their consistent derivation from in situ observations remains challenging because size distribution, absorption, and humidity effects are coupled. Here we present a decoupled inversion framework that integrates SMPS–OPC size distributions, AE33-derived absorption coefficients, and ambient relative humidity observations from the Vipava Valley (Slovenia) in April 2016 to derive dry-state homogeneous-equivalent CRI and RH-corrected lidar-

relevant SSA and LR at 355, 532, and 1064 nm. The framework reconstructs a continuous particle size distribution through instrument-response-aware SMPS–OPC geometric alignment and retrieves the reference real refractive index. AE33-derived absorption constraints and a singly subtractive Kramers–Kronig (SSKK) relation are then used to derive the dry-state CRI spectrum, after which SSA and LR are recalculated through an RH-corrected wet-state forward optical calculation. Under the assumed AE33 reference scenario with an effective multiple-scattering correction factor of C_{eff} = 5, mean dry-state n values were 1.449, 1.441, and 1.437, and mean k values were 0.0309, 0.0235, and 0.0190 at 355, 532, and 1064 nm, respectively. The RH-corrected mean SSA values were 0.855, 0.841, and 0.760, and the corresponding mean LR values were 86.53, 57.59, and 39.81 sr. Increasing C_{eff} from 1.57 to 5 decreased median RH-corrected LR_{355} from 137.3 to 84.3 sr and increased median SSA_{355} from 0.646 to 0.867. These conditional, site-specific estimates link dry in situ retrievals with RH-corrected lidar-relevant optical properties.

Comment 2: I appreciate the additional objective-profile and structural-sensitivity analyses provided in Response 4.4, which improve the assessment of the SMPS-OPC retrieval. However, Figure R3(a) indicates variations in the identifiability of (n_{655}), including cases with objective-function support extending to the LUT boundaries. Since these diagnostics were not used as retrieval-screening criteria, apart from excluding five solutions with (n_{655} < 1.33), could the authors clarify how many of the final 304 retrievals exhibit broad or boundary-reaching objective-function support and whether these cases materially influence the reported (n_{655}) statistics? A brief summary of the existing diagnostic results would be sufficient.

Response: We thank the referee for this suggestion. Using the descriptive relative-5% support defined in Supplement Sect. S2, 17 of the 304 profiles (5.6%) reach a LUT boundary (8 lower, 9 upper). Their best-fit values remain internal; support beyond the LUT was not assessed. Using $W_{5\%} \geq 0.1$, one-quarter of the 1.30–1.70 LUT span, as a descriptive width reference, 84 profiles (27.6%) have wider support. Sixteen also reach a boundary, giving 85 profiles (28.0%) in the combined group. This width reference is not a validated screening threshold. Descriptive comparisons of the n_{655} summary statistics are presented in Table R1.

Table R1. Mean and median n_{655} for the full cohort and diagnostic subsets.

Summary subset	***N***	**Mean n_{655}**	**Median n_{655}**
Full final cohort	304	1.4399	1.4155
Support contained within the LUT	287	1.4377	1.4124
$W_{5\%} < 0.10$	220	1.4205	1.3886

Omitting boundary-contact cases from the statistical summary changes the mean and median by −0.0022 and −0.0031, respectively; restricting to $W_{5\%} < 0.1$ produces larger changes of −0.0193 and −0.0269. These differences do not estimate retrieval bias or establish greater accuracy for narrower-support cases. All 304 retrievals remain in the analysis, with no additional screening criterion introduced. The diagnostic definitions and numerical summaries have been added to Supplement Sect. S2.

Comment 3: The PSD-truncation analysis provided in Response 5.1 is useful and clarifies the contribution of the coarse tail to LR_{1064}. I note, however, that the retrieved CRI was held fixed while the PSD was truncated. The experiment therefore quantifies the direct forward-optical sensitivity to the assumed upper-size cutoff, rather than the combined sensitivity of the full inversion, in which PSD changes could also affect the retrieved refractive index. Please make this distinction explicit when interpreting the PSD-truncation results. No additional full retrieval experiment is necessarily required.

Response 3: We agree. The truncation tests were designed to isolate the direct forward-optical effect of changing the PSD upper-size cutoff. We therefore held fixed the dry-state CRI obtained for each hour and wavelength from the original full-PSD retrieval, without refitting *n* or *k*.

These tests changed the dry-PSD upper-size cutoff while holding the retrieved CRI, RH, hygroscopicity parameters, and wet-state mixing treatment fixed; retained concentrations were not renormalized. The resulting changes therefore quantify direct forward-optical sensitivity to the cutoff, rather than the combined sensitivity of the full inversion. We have made this distinction explicit in Sect. 4.3, consistent with the experimental design already described in Supplement Sect. S5.

Comment 4: Response 4.3 clarifies that the AE33 sampled through a $PM_{2.5}$ inlet, whereas the Mie-based absorption closure appears to use the merged PSD extending to 20 um. Although (k) is retrieved as a homogeneous-equivalent parameter constrained by bulk absorption, this difference in particle-size coverage may influence its interpretation and the subsequent optical calculations. I suggest briefly acknowledging this limitation and clarifying its potential implications for the retrieved (k), SSA, and LR in Section 4.3.

Response: We agree. The AE33 absorption measurements were obtained downstream of a $PM_{2.5}$ inlet, whereas the Mie absorption closure uses the merged PSD extending to approximately 20 µm. The size-coverage mismatch may influence the fitted homogeneous-equivalent *k* and propagate into SSA and LR calculated for the full PSD. We have acknowledged this limitation in Sect. 4.3, including the distinction between the aerodynamic inlet cutoff and the reconstructed PSD diameter coordinate. The available observations do not establish the magnitude or direction of the resulting bias.

Changes in the manuscript — Sect. 4.3, line 556–562

The AE33 absorption measurements were obtained downstream of a $PM_{2.5}$ inlet, whereas the Mie absorption closure uses the merged PSD extending to approximately 20 µm. The aerodynamic inlet cutoff is not equivalent to a sharp 2.5 µm cutoff on the reconstructed PSD diameter coordinate. Fitting full-PSD model absorption to this inlet-limited measurement may cause the homogeneous-equivalent *k* to partly reflect the size-coverage mismatch. Applying this fitted *k* across the full PSD may misrepresent coarse-particle absorption and affect the derived SSA and LR; coarse-particle absorption is not independently constrained. The fixed-CRI PSD-truncation tests do not assess this potential response of the fitted *k*. The available observations do not establish the magnitude or direction of the resulting bias.