Response to the Comments of Referees

We would like to thank to the reviewers for giving constructive criticisms, which are very helpful in improving the quality of the manuscript. We have made major revision based on the critical comments and suggestions of the referees. The referee's comments are reproduced (*black*) along with our replies (blue) and changes made to the text (red) in the revised manuscript. All the authors have read the revised manuscript and agreed with submission in its revised form.

Anonymous Referee #2

Comment NO.1: *This manuscript presents a component-resolved 4D-LETKF data assimilation system that assimilates CALIOP-MODIS synergistic speciated extinction retrievals (sea salt, dust, BC, water- soluble) into WRF-Chem/GOCART over East Asia, with the explicit goal of correcting the species- misallocation problem inherent to traditional total-AOT or bulk-extinction assimilation. The system is evaluated against the assimilated CALIOP profiles themselves, and independently against AERONET AOT/SSA and AD-Net lidar extinction.*

The headline result is the reconstruction of a persistent free-tropospheric BC layer over the Indian subcontinent that the free-running model misses entirely, along with corresponding corrections to single scattering albedo and shortwave heating rate profiles. The problem motivation is sound, the observational dataset is valuable, and the independent AERONET/SSA validation is a genuine strength.

Several aspects of the assimilation design complicate the interpretation of the reported skill. Specifically, the background-correction bypass mechanism, the construction of the ensemble, and the combined water-soluble observation operator make it difficult to distinguish LETKF flow-dependent covariance skill from simpler bias-rescaling. These concerns must be resolved before the central claims can be fully assessed. I recommend major revisions.

Response: We sincerely thank the reviewer for the highly professional and constructive evaluation of our manuscript. We are greatly encouraged by the positive recognition of our problem motivation, the observational dataset, and the independent AERONET validation.

We fully understand your overarching concerns regarding whether the reported skill stems from LETKF flow-dependent covariance or from the assimilation design (specifically the bypass mechanism, the ensemble construction, and the observation operator). To comprehensively resolve these problems, we have conducted a new sensitivity assimilation experiment (DA_NoBypass) and restructured the manuscript by adding a new Section 5.2 (Discussion on the DA system limitations and the bypass mechanism) to explicitly discuss the system's limitations and true capabilities.

Detailed, point-by-point responses regarding the bypass mechanism, the construction of the ensemble, and the combined water-soluble observation operator are provided in our specific responses to Comments 2, 3, and 4 below.

Comment NO.2: *The background-correction scheme (Section 3.3) rescales the background ensemble by 1/R_corr when the localized weighted bias ratio falls below 0.2, bypassing the LETKF* *entirely. Because R_corr is constructed from the ratio of background to observations, this step drives the analysis toward the observations wherever the bypass triggers. This raises a critical concern: the near-perfect domain-mean correlations reported in Figure 4 (CORR = 1.00 for dust and water- soluble; CORR = 0.95 for BC) may partially reflect this forced agreement rather than genuine ensemble-based state-estimation skill.*

The issue is most acute for black carbon, where the FR simulation completely misses the elevated 6 km peak (Figure 4c; CORR = 0.26) and background ensemble spread collapses to near- zero (lines 283-285). Under these conditions, the standard LETKF would yield negligible Kalman gain and effectively ignore the observations. The bypass therefore likely dominates the update for these missed plumes, and the resulting DA CORR of 0.95 is achieved primarily by rescaling rather than by flow-dependent covariance. For dust and water-soluble aerosols, where the FR simulation already captures the general vertical structure (CORR = 0.98 and 0.90 respectively), the marginal improvement to 1.00 is similarly difficult to attribute to LETKF performance without knowing how often the bypass intervened.

To resolve this ambiguity and allow readers to evaluate the true LETKF contribution, I suggest the paper report:

• The fraction of grid points, vertical levels, and assimilation cycles where the bypass triggered versus standard LETKF assimilation, disaggregated by species;

• The Figure 4 statistics computed separately for bypass-corrected and LETKF- assimilated subsets, so readers can evaluate whether the dramatic CORR/RMSE improvements reflect ensemble covariance skill or bias-rescaling;

• Ideally, a sensitivity experiment omitting the bypass (or with a more stringent threshold) to isolate the LETKF's contribution under conditions of near-zero background spread.

Without this decomposition, it is difficult to distinguish LETKF flow-dependent covariance skill from simpler bias-rescaling, and the paper's central claim that component-resolved ensemble assimilation improves vertical structure remains incompletely supported.

Response: We deeply appreciate this rigorous and highly constructive suggestion. We fully agree that decomposing the contribution of the bypass mechanism is essential for evaluating the true LETKF performance. To demonstrate this, we have conducted a new assimilation experiment (DA_NoBypass) as pure assimilation. We have addressed all three suggestions as follows:

(1) Triggering Statistics: We have calculated the exact triggering frequency. Spatially and temporally, the bypass mechanism was triggered in approximately 28.6% of the total assimilated grid points across all assimilation cycles. Disaggregated by species, it triggered in 43.6% of cases for dust, 17.8% for water-soluble aerosols, and 24.5% for BC. Vertically, it predominantly occurred in the free troposphere (e.g., 4-7 km) rather than the boundary layer.

(2) Separated Statistics: As shown in Fig. 16, the DA_NoBypass experiment significantly improves the CORR across different aerosol species compared to the FR experiment. For instance, the CORR for Black Carbon (BC) dramatically improved from 0.26 (FR) to 0.79 (DA_NoBypass).

Similar structural improvements are also observed in sea-salt (CORR improved from 0.87 to 0.93) and water-soluble (from 0.90 to 0.99). Because simple bias-rescaling primarily corrects systematic mean errors without significantly altering the temporal or spatial variation patterns, these substantial improvements in CORR particularly for BC, strongly demonstrate that the LETKF system possesses genuine flow-dependent ensemble covariance skill. It effectively captures and corrects the dynamic distribution of aerosols, rather than just performing a numerical rescaling. Therefore, the extreme local improvements in completely unsimulated plumes highlight a fundamental limitation of ensemble-based DA: empirical corrections remain indispensable when background spread collapses. Conversely, the independent AERONET and AD-Net validations presented in Section 4.2 provide a much more robust and conservative estimate of the system's true predictive skill. Because these sites are often far from CALIOP overpass tracks, the modest but solid improvements observed there reflect the genuine dynamical propagation of the assimilated fields rather than the forced local agreement of the bypass mechanism.

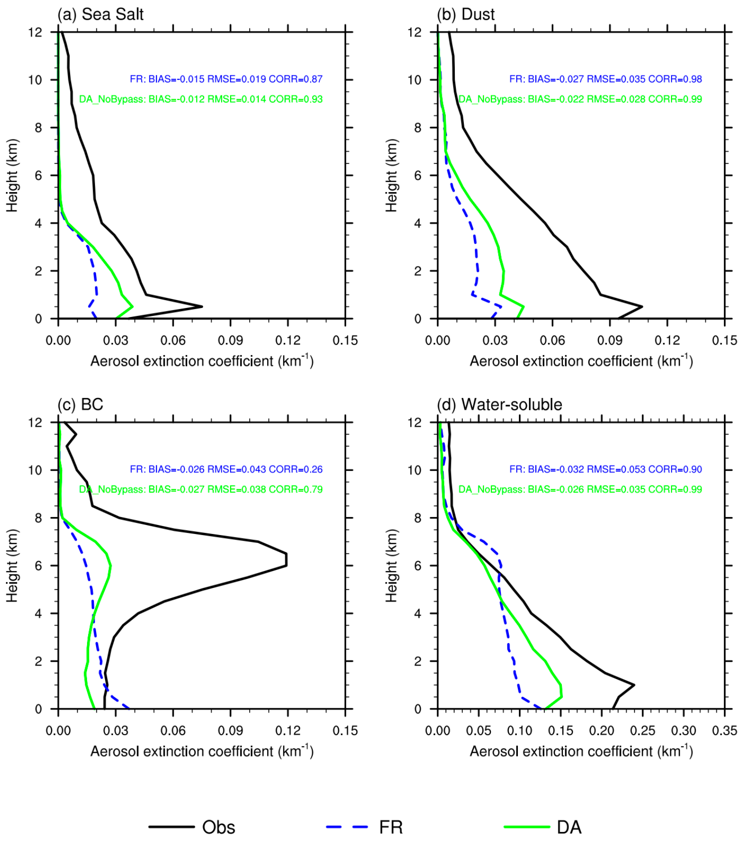

Figure 16. Same as Fig. 4 but for DA_NoBypass experiment.

(3) Sensitivity Experiment: As recommended, we conducted a new experiment (DA_NoBypass) omitting the bypass mechanism entirely. The results show that even under conditions of near-zero background spread, the DA_NoBypass experiment still yields substantial improvements (e.g., sea salt, dust, black carbon, and water-soluble aerosol CORR improving to 0.93, 0.99, 0.79, and 0.99), though it naturally avoids the unrealistic over-fitting seen in the DA experiment.

All these quantitative decompositions and the new sensitivity experiment have been thoroughly documented in the newly added Section 5.2.

Changes in Manuscript: Please refer the description in new Section 5.2 in the revised manuscript as follows:

"5.2 Discussion on the DA System Limitations and the Bypass Mechanism

While the component-resolved assimilation system demonstrates significant improvements, it is crucial to objectively distinguish the contributions of the background correction mechanism from the pure LETKF assimilation. It should be acknowledged that the background correction scheme acts as a simplified nudging approach that bypasses the rigorous uncertainty weighting of the standard Kalman filter. Consequently, this direct rescaling becomes less trustworthy when observational uncertainties are large. Although we implemented a 3-sigma outlier check and required a minimum number of local observations to mitigate the impact of noisy retrievals, ignoring observational uncertainty during the correction step remains a limitation.

To demonstrate the true capability of the LETKF, we have conducted a new assimilation experiment (DA_NoBypass) as pure assimilation. Spatially and temporally, the bypass mechanism was triggered in approximately 28.6% of the total assimilated grid points across all assimilation cycles. Disaggregated by species, it triggered in 43.6% of cases for dust, 17.8% for water-soluble aerosols, and 24.5% for BC. Vertically, it predominantly occurred in the free troposphere (e.g., 4-7 km) rather than the boundary layer. As shown in Fig. 16, the DA_NoBypass experiment significantly improves the CORR across different aerosol species compared to the FR experiment. For instance, the CORR for Black Carbon (BC) dramatically improved from 0.26 (FR) to 0.79 (DA_NoBypass). Similar structural improvements are also observed in sea-salt (CORR improved from 0.87 to 0.93) and water-soluble (from 0.90 to 0.99). Because simple bias-rescaling primarily corrects systematic mean errors without significantly altering the temporal or spatial variation patterns, these substantial improvements in CORR particularly for BC, strongly demonstrate that the LETKF system possesses genuine flow-dependent ensemble covariance skill. It effectively captures and corrects the dynamic distribution of aerosols, rather than just performing a numerical rescaling.

Therefore, the extreme local improvements in completely unsimulated plumes highlight a fundamental limitation of ensemble-based DA: empirical corrections remain indispensable when background spread collapses. Conversely, the independent AERONET and AD-Net validations presented in Section 4.2 provide a much more robust and conservative estimate of the system's true predictive skill. Because these sites are often far from CALIOP overpass tracks, the modest but solid improvements observed there reflect the genuine dynamical propagation of the assimilated fields rather than the forced local agreement of the bypass mechanism."

Comment NO.3: *The 20-member ensemble is generated solely by perturbing aerosol emissions(Section 3.3). So the meteorology is running deterministically? This means the ensemble has negligible spread in vertical transport processes (convective lofting, PBL mixing, plume-rise injection heights). But in line 321-323, the authors attribute the persistent underestimation of the 6* *km BC peak magnitude to "deficiencies in the model's vertical transport mechanisms or an underrepresentation of the biomass burning plume injection heights".*

This creates an internal tension: if the background covariance contains no information about transport uncertainty, the system can only redistribute mass vertically via the prescribed emission profiles, not correct the altitude of lofted plumes through flow-dependent covariance. The fact that DA captures the shape but not the magnitude of the 6 km BC peak (Figure 4c) is consistent with this structural limitation. The authors should explicitly address this constraint.

Response: Agree. Because the meteorological fields are run deterministically, the ensemble spread is generated exclusively by perturbing the aerosol emissions. Consequently, the background error covariance matrix primarily reflects uncertainties in emission magnitudes and lacks flow-dependent uncertainties related to vertical transport processes. As a result, when the forward model fails to loft the BC plumes to the 6 km altitude due to deficiencies in vertical transport mechanisms, all ensemble members simulate near-zero concentrations at this level. This leads to a negligible ensemble spread at 6 km. Under such conditions, the standard LETKF struggles to redistribute mass vertically because the flow-dependent covariance cannot correct the altitude of the lofted plumes. Although our background correction strategy helps capture the general shape of the elevated layer, it cannot fully compensate for the lack of transport-driven spread, which fundamentally restricts the system's ability to reproduce the absolute magnitude of the 6 km BC peak.

Changes in Manuscript: We have revised Section 4.1 in the manuscript to clearly state this structural limitation and its consequences on the assimilated vertical profiles:

"Because our assimilation system generates ensemble spread solely by perturbing aerosol emissions while running the meteorology deterministically, the background error covariance lacks flow-dependent information regarding vertical transport uncertainties. Consequently, when the model fails to loft BC to the 6 km altitude, the ensemble spread at this level approaches zero. The DA system can only redistribute mass based on the prescribed emission-driven covariance, limiting its ability to fully correct the altitude of lofted plumes. This limitation explains why the DA experiment captures the shape of the 6 km BC peak but still underestimates its absolute magnitude (Fig. 4c)."

Please refer the description in the revised manuscript, Line 411-417.

Comment NO.4: *The paper acknowledges that GOCART assumes externally mixed BC, whereas the CALIOP-MODIS retrieval assumes internal (core-shell) mixing (Section 4.1, lines 288-292). However, this discrepancy is not carried through to the radiative heating evaluation in Section 4.4. Since the heating rates are computed using the model's own (externally mixed) optical scheme applied to the DA-analysis BC mass, any bias in assumed mass absorption efficiency propagates directly into the reported 3.0 K/day peak heating anomaly. I would ask the authors to comment, even qualitatively, on the expected sign and approximate magnitude of this effect, since it bears directly on the paper's quantitative radiative. Specifically, the authors should shed some light regarding:*

- *Is the DA system inflating BC mass beyond its true atmospheric value to match observed extinction;*

• *Is the diagnosed SW heating rates (Figure 13) then applying this excess mass to the model's lower MAC, creating a non-obvious bias in the 3.0 K day^{-1} peak heating estimate.*

The authors should comment on the sign and approximate magnitude of this effect. Is the reported peak heating an upper bound, a lower bound, or bracketed by the two mixing assumptions? Is the reported peak heating an upper bound, a lower bound, or bracketed by the two mixing assumptions? Given recent literature cited in the Discussion (Tiwari et al., 2025; Guan et al., 2026a), this could perturb regional forcing estimates by several W m^{-2}.

Response: Done.

(1) Yes. The CALIOP-MODIS retrieval assumes an internal core-shell mixing state, which typically exhibits a higher mass extinction efficiency (MEE) due to the lensing effect. In contrast, the GOCART model assumes externally mixed BC, which has a lower MEE. To match the assimilated extinction observations using a lower MEE, the DA system inevitably has to inflate the BC mass beyond its true atmospheric value.

(2) The reported 3.0 K day^{-1} peak heating is likely a lower bound. The shortwave heating rate is fundamentally driven by the absorption coefficient (β_{abs}). In our DA system, the analyzed extinction matches the observation ($\beta_{ext,DA} \approx \beta_{ext,obs}$). The model calculates absorption as $\beta_{abs,DA} = \beta_{ext,obs} \times (1 - SSA_{ext})$. However, the true absorption would be $\beta_{abs,true} = \beta_{ext,obs} \times (1 - SSA_{int})$. Because internal mixing significantly enhances absorption, its single scattering albedo (SSA) is lower than that of external mixing ($SSA_{int} < SSA_{ext}$). Therefore, the model's assumed absorption fraction per unit extinction is lower than in reality. Despite the inflated BC mass, the underestimation of the mass absorption cross-section dominates, meaning the actual atmospheric heating induced by this BC plume is likely even stronger than our 3.0 K day^{-1} estimate.

Changes in Manuscript: We have added a qualitative discussion of this effect in Section 4.4 to clarify that our estimate serves as a conservative lower bound:

"It is important to note that the estimated 3.0 K day^{-1} peak heating likely represents a conservative lower bound due to discrepancies in aerosol mixing state assumptions. The CALIOP-MODIS retrieval assumes an internal core-shell mixing state, whereas the GOCART optical scheme assumes externally mixed BC. Because external mixing yields a lower mass extinction efficiency, the DA system must inflate the BC mass to match the observed extinction. However, internal mixing significantly enhances absorption, resulting in a lower single scattering albedo. Consequently, while the DA system successfully constrains the total extinction, applying the externally mixed optical scheme underestimates the absorption fraction per unit extinction. Therefore, the true shortwave heating induced by these elevated BC plumes could be even stronger than our current estimates."

Please refer the description in the revised manuscript, Line 568-575.

Comment NO.5: *The observation operator for water-soluble (WS) extinction combines sulfate and organic carbon (OC) into a single channel: β_WS = α_OC·M_OC + α_SU·M_SU (Section 3.1). While the paper criticizes traditional total-AOT assimilation for being under-constrained with* *respect to aerosol composition, this WS channel creates a similar problem at the species level: a single optical observation cannot independently constrain two chemically and radiatively distinct species. The LETKF partitions the increment using background ensemble cross-correlations between OC and sulfate, but these species have different sources, lifetimes, hygroscopic growth, and deposition behaviors. The variable localization (Section 3.2) permits WS observations to update both species, but the paper provides no validation of the resulting OC/sulfate partition. Given that both OC and sulfate are predominantly scattering-dominated aerosols, mispartitioning between them likely has much smaller radiative consequences than misallocating between BC and scattering species. However, from the perspective of aerosol composition and chemical transport, the OC/sulfate partition is still a relevant uncertainty. Additionally, if future applications of this system aim to use the analysis fields for process studies or source attribution, the inability to independently constrain OC and sulfate would be a limitation. The authors should acknowledge this under-constrained aspect of the WS observation operator explicitly as a limitation and clarify whether the LETKF's ensemble correlations between OC and sulfate are physically reasonable or whether the system tends to adjust one species preferentially.*

Response: We sincerely thank the reviewer for this highly insightful comment. We completely agree that combining organic carbon (OC) and sulfate (SU) into a single water-soluble (WS) extinction channel creates an under-constrained problem at the species level, as a single optical observation cannot independently constrain two chemically and radiatively distinct species.

Following your constructive suggestion, we have explicitly acknowledged this limitation and clarified the data assimilation mechanisms regarding the OC/sulfate partitioning in Section 3.1 of the revised manuscript. Specifically, we address your queries regarding as follows:

(1) Physical reasonableness of ensemble correlations: Because our 20-member ensemble is generated solely by perturbing aerosol emissions under deterministic meteorology, the background cross-correlations between OC and sulfate in the LETKF are primarily driven by the spatial co-location of their emission sources and their subsequent co-transport. While these correlations are physically reasonable for capturing the macro-structure of co-emitted aerosol plumes, we acknowledge that they cannot fully capture the slight differences in aerosol processes between the two species.

(2) Preferential adjustment by the LETKF: The LETKF modify the analysis increment based on the background error covariance. Consequently, the system does not have a predefined physical preference; rather, it tends to preferentially adjust the species that exhibits a larger background variance and a higher mass extinction efficiency in the observation operator.

(3) Limitations of the observation operator: It is important to acknowledge that combining OC and sulfate into a single water-soluble extinction (β_{WS}) creates an under-constrained problem at the species level. Since a single optical observation cannot separate two chemically distinct species, the LETKF partitions the analysis increments relying entirely on background ensemble cross-correlations.

Changes in Manuscript: We have added the description in the Section 3.2 as below:
"However, it is important to acknowledge that combining OC and sulfate into a single water-soluble extinction (β_{WS}) introduces an inherent limitation, as it creates an under-constrained problem at the species level. Since a single optical observation cannot separate two chemically distinct species, the LETKF must partition the analysis increments relying entirely on background ensemble cross-correlations. Because our 20-member ensemble is generated solely by perturbing aerosol emissions under deterministic meteorology, the background cross-correlations between OC and sulfate are primarily driven by the spatial co-location of their emission sources and their subsequent co-transport. While these correlations are physically reasonable for capturing the macro-structure of co-emitted aerosol plumes, we acknowledge that they cannot fully capture the slight differences in aerosol processes between the two species. Consequently, the LETKF modifies the analysis increment based on the background error covariance without a predefined physical preference; rather, the system tends to preferentially adjust the species that exhibits a larger background variance and a higher mass extinction efficiency in the observation operator." Please refer the description in the revised manuscript, Line 237-247.

Comment NO.6: *The DA experiment reports near-perfect domain-mean correlations against the assimilated CALIOP extinction profiles in Figure 4 (CORR = 1.00 for dust and water-soluble; CORR = 0.95 for BC), representing dramatic improvements over the FR experiment (0.98, 0.90, and 0.26, respectively). However, these self-validation metrics stand in striking contrast to the independent AERONET validation in Figure 6, where the column-integrated AOT shows only modest improvement. CORR increases from 0.573 to 0.626, BIAS improves from -0.278 to -0.208, and the fraction of samples within ±0.10 of the observed AOT rises only from 23.36 % to 28.10 %. While it is expected that self-validation against assimilated data would be stronger than independent validation, the magnitude of the discrepancy raises questions about how much of the self-validation skill reflects genuine state-estimation capability versus forced agreement with the assimilated observations. The near-perfect correlations (1.00 for two species) are unusually high for a 20-member ensemble system with sparse satellite observations and could indicate over-fitting or the influence of the bypass mechanism described in Major Comment 1.*

This concern is amplified by the background-correction bypass mechanism (Section 3.3). When the background ensemble spread collapses to near-zero-which occurs precisely for the elevated BC plumes at ~6 km that dominate the FR error (Figure 4c, lines 283-285) (also mentioned in Major Comment 1), the system rescales the background by 1/R_corr and bypasses the LETKF entirely. The resulting CORR = 0.95 for BC may therefore be partially tautologically guaranteed wherever the bypass dominates, rather than earned through flow-dependent covariance. The AERONET validation, being independent and located at sites often far from CALIOP overpass tracks, is not subject to this forced agreement, which may explain why the skill improvement there is more modest. The spatial distribution of AERONET improvements (Figure 7) further illustrates this point. While the DA experiment reduces absolute BIAS at 91.5% of sites and RMSE at 71.2%, the magnitudes *are heterogeneous, and some sites show negligible or even negative improvement. Given that CALIOP tracks are sparse and the horizontal localization radius is 75 km, improvements at downwind sites necessarily come from dynamical propagation of the analysis state rather than direct observation influence. The patchy nature of the improvements is therefore not inherently problematic, it reflects the sparsity of the satellite observations and the complexity of transport pathways.*

Nevertheless, the contrast between near-perfect CALIOP correlations and modest AERONET improvements requires explicit reconciliation. I therefore ask the authors to:

1. Address whether the CALIOP CORR values of 1.00 are physically credible or may reflect over-fitting;

2. Discuss the role of the bypass mechanism in generating these near-perfect correlations;

3. Acknowledge that the independent validation provides a more conservative estimate of the system's true performance.

Without this reconciliation, the manuscript presents two seemingly incompatible pictures of performance, and resolving which picture reflects the true capability of the component-resolved LETKF is essential before the central claims can be fully assessed.

Response: We concur entirely with the reviewer's assessment and have explicitly addressed all three points in the revised manuscript:

(1) Over-fitting: We explicitly acknowledge that the CALIOP CORR values of 1.00 are physically unrealistic for a 20-member ensemble and represent mathematical over-fitting.

(2) Role of bypass: This tautological agreement is a direct consequence of the bypass mechanism, which forces the analysis state to match the assimilated observation when the background spread collapses. Our new DA_NoBypass experiment confirms that without this mechanism, the pure LETKF yields more physically credible, although lower, correlations.

(3) Conservative independent validation: We fully agree that the independent AERONET and AD-Net validations provide a much more conservative and realistic metric of the system's true predictive skill. Because these sites are often far from CALIOP tracks, the improvements there are earned through genuine dynamical propagation rather than forced local agreement.

Changes in Manuscript: Please refer the description in new Section 5.2 in the revised manuscript.

Comment NO.7: *If the background ensemble mean is effectively zero at a grid point (a completely missed aerosol event), R_corr approaches zero and the correction step requires division by a near-singular value. The paper mentions a 3-sigma outlier check (line 228), but does not describe an explicit numerical floor or regularization for R_corr. Given that missed BC plumes are the primary motivation for the correction scheme, the authors should clarify how this edge case is handled (e.g., is R_corr floored at 0.01? Is a minimum background spread enforced?).*

Response: Agree. A completely missed aerosol event could lead to a near-singular division. To address this, explicit numerical constraints are implemented in the assimilation algorithm to handle such edge cases. Specifically, two thresholds are applied to the scaling factor R_corr before the correction step:

(1) A minimum threshold of 1.0×10^{-15} is required to prevent division by zero.

(2) During the 3-sigma quality control step, the lower bound of the valid R_{corr} range is explicitly set to 0.001.

Consequently, the background correction is only executed when $R_{corr} > 0.001$. For grid points where the background aerosol concentrations are effectively zero (resulting in $R_{corr} < 0.001$), the correction step is skipped. This constraint effectively prevents numerical instability and spurious inflation of the background state.

Changes in Manuscript: We have added the description in the Section 3.4 as below:

"To maintain numerical stability when the background ensemble mean is close to zero (e.g., in the case of completely missed aerosol plumes), two thresholds are applied to the scaling factor R_corr before the correction step: (1) a minimum threshold of 1.0×10^{-15} is set to prevent division by zero; (2) during the 3-sigma outlier check, the lower bound for valid R_{corr} values is constrained to 0.001. If the calculated R_{corr} is less than 0.001, the background correction is not applied at that specific grid point. This constraint effectively prevents spurious inflation of the aerosol state when the background concentrations are negligible." Please refer the description in the revised manuscript, Line 304-309.

"Comment NO.8: *The simulation uses HTAPv3 emissions for 2018, while the study period is March 2021- a three-year gap that includes the COVID-19 emission perturbation. The authors should acknowledge that some portion of the diagnosed "model structural deficiency" may reflect stale emission inputs rather than transport or mixing errors.*

Response: Agree. The use of the 2018 emission inventory for a 2021 simulation inevitably introduces biases, particularly given the emission changes during the COVID-19 pandemic. We have added a statement in the revised manuscript to clarify this.

Changes in Manuscript: We have added the description as below:

"It should be noted that a portion of the diagnosed model structural deficiencies likely stems from the use of the 2018 HTAPv3 emission inventory for the 2021 study period, which does not account for emission perturbations associated with the COVID-19 pandemic." Please refer the description in the revised manuscript, Line 121-123.

Comment NO.9: *The temporal localization scale σ_t = 0.01 is given without units. Please clarify what this value represents physically.*

Response: The temporal localization scale $\Delta t = 0.01$ is in units of time slots (where one slot corresponds to the temporal interval of the assimilation window). Physically, this scale acts as a strict temporal filter. Specifically, setting this value to 0.01 means that only observations at the exact current time step are utilized for the assimilation update, effectively excluding any asynchronous observations within the assimilation window.

Changes in Manuscript: Please refer the description in the revised manuscript, Line 274-278.

Comment NO.10: *Tsikerdekis et al. (2021), already cited, also performed speciated optical-property assimilation (POLDER/PARASOL size and absorption constraints). The introduction should more explicitly state what distinguishes the present approach (e.g., explicit per-species vertical mapping, the background-correction step) from that prior work.*

Response: Accept. We have added the description as below:

"While previous studies, such as Tsikerdekis et al. (2021), have successfully performed speciated aerosol assimilation using column-integrated, multi-angle, and polarization measurements to constrain aerosol size and absorption, our approach leverages high-resolution vertical extinction profiles from CALIOP to provide explicit vertical mapping of different aerosol components."

Changes in Manuscript: Please refer the description in the revised manuscript, Line 93-96.

Comment NO.11: *In Figure 15, the latitude-height heating-rate cross-section is shown for a single timestamp (06:00 UTC, 7 March 2021), but the accompanying text generalizes the "warming aloft, reduced heating below" structure as a regional feature!!!! Either a composite over multiple events or a more explicit caveat that this is a single case study would be appropriate should be mentioned.*

Response: Agree. We have replaced "To investigate the physical mechanisms of aerosol radiative effects across different source regions, Figure 15 presents the latitude-height cross-sections of speciated aerosol extinction and the corresponding shortwave (SW) heating rates" with "To investigate the physical mechanisms of aerosol radiative effects across different source regions, Figure 15 presents a specific case study at 06:00 UTC on March 7, 2021, illustrating the latitude-height cross-sections of speciated aerosol extinction and the corresponding shortwave (SW) heating rates."

Changes in Manuscript: Please refer the description in the revised manuscript, Line 593-595.