We thank the reviewers and community commenter for their careful assessments and constructive comments. Their suggestions helped us clarify the physical basis for the hybrid BNN training set and the advantages hybrid BNN framework, strengthen the uncertainty and sensitivity analyses, expand comparisons with established optical approximations for fractal aggregates, and improve the documentation and reproducibility of the ADDA simulations. Our detailed responses are provided below, with reviewer comments reproduced in black and our responses in blue.

Reviewer 1:

Assessment:

This study introduces a hybrid Bayesian neural network (BNN) model trained on discrete dipole approximation (DDA) simulations for black carbon (BC)-containing particles. The description of the proposed model is generally clear. A total of eight input features are considered: five particle-related inputs (number of BC particles, coating amount, fractal dimension, volume-equivalent size parameter, and imaginary refractive index of the coating) and three additional inputs derived from homogeneous-sphere Mie calculations (extinction coefficient, single-scattering albedo, and asymmetry parameter). The three model outputs consist of optical properties (extinction coefficient, single-scattering albedo, and asymmetry parameter) from DDA simulations. In addition, the treatment and decomposition of aleatoric and epistemic uncertainty are generally well explained.

Overall, the manuscript is well written, well organized, and easy to follow, except for some details in the machine-learning-related sections where I am not an expert. The scope is well addressed, and the general methodology and evaluation approaches appear reasonable. However, several aspects would benefit from additional clarification. Therefore, I recommend publication after the authors address the comments listed below.

Specific Comments:

1. Different assumptions between Ensemble I and II (Table 1 and related discussion): Why do differences exist with respect to wavelength, coating refractive index, AAE of the coating, and other parameters?

   The differences between Ensembles I and II reflect differences in the underlying physical model that was used to generate the BC core. Ensemble I was generated before a physical BC restructuring model was available, and was used to develop scaling laws for BC absorption as a function of phase shift parameter, coating amount, wavelength, and coating absorptivity (Beeler and Chakrabarty, 2022). Discrete values for wavelength, coating absorptivity, etc. were used to aid in fitting of the scaling laws developed in Beeler and Chakrabarty (2022). These simulations were used in this work to represent edge cases for BC shape, coating amount, and wavelength-dependence of coating refractive index.

   On the other hand, Ensemble II was generated using an observationally-constrained BC compaction model (Beeler et al., 2025), where the fractal dimension of the BC core increases as a function of coating amount. The ADDA simulations of these particles were performed for the purposes of training the BNN model, and therefore we varied the input parameters in a more systematic fashion. The advantage of including both ensembles in the training/validation datasets is that the effects of each input feature on the predicted optical properties can be discerned because the input features have differing correlations with one another between the two ensembles. We have added more information on the distinction between Ensembles I and II in Sections 2.1 and 2.2 of the revised manuscript.

2. Figure 1: What is the distinction between the blue and brown particles? I assume the partially transparent regions indicate coatings, but it was not entirely clear whether the figure simply represents uncoated versus coated particles.

   The reviewer is correct that the partially transparent regions indicate coating material, while the blue represents the BC core. We have added a legend to the revised Figure 1 to clarify this distinction.

3. How is the core fractal dimension ($D_{f,c}$) calculated? It would be helpful to provide at least a brief definition or formulation in the Appendix. In addition, the definition or relationship of $D_{f,c}$ with other parameters appears to differ between Ensembles I and II, which was somewhat difficult to follow.

   We now clarify in lines 81-84 and 101-104 of the revised text that the core fractal dimension of each particle in the training ensembles is calculated using the structure factor approach of Sorensen (2001). This calculation is the same for both ensembles, which we now make clear in lines 81-84 and 101-104 of the revised text; what differs between the two ensembles is how the particles are generated. In Ensemble I, core morphology and coating amount were varied independently, without accounting for core restructuring as particles are coated; in Ensemble II, the particle-generation procedure couples the core morphology to V/V_0 and $N_p p$ using Equations 1–3, which were derived from laboratory measurements and detailed models of coating-induced capillary compaction, as described in Beeler et al. (2025). We have revised Figure 1 to show that $D_{f,c}$ and V/V_0 are statistically independent in Ensemble I, denoted by $D_{f,c} \perp\!\!\!\perp V/V_0$. We also show that $D_{f,c}$ and V/V_0 are not statistically independent in Ensemble II, denoted by $D_{f,c} \not\perp\!\!\!\perp V/V_0$.

4. L79: Is it reasonable to assume a single fixed value for the BC refractive index? Does the BC refractive index have a known uncertainty range or variability that should be considered?

   The reviewer is right that the refractive index of BC has been found to vary from less refractive ($1.7+0.64i$) to highly refractive ($2.0 + 1.0i$). In Section 7 of the original manuscript, we evaluated the sensitivity of BNN predictions to different assumed values of the BC refractive index (panels B, E, H of Figure 6). We show that prediction errors for Q_{ext} increased when the BC refractive index departed from the value used in the training set ($1.95 + 0.79i$), but the error produced by the hybrid BNN remained comparable to those from core-shell and homogeneous sphere approximations. We also found that the hybrid BNN is more accurate than core-shell and homogeneous sphere approximations for SSA and g across the tested BC refractive indices. This analysis indicates that the homogeneous-sphere optical descriptors allow the hybrid BNN to retain useful predictive skill for BC refractive indices not explicitly represented in the training set, provided that $Q_{ext,\mathrm{hs}}$, $\mathrm{SSA}_{\mathrm{hs}}$, and g_{hs} are recalculated using the desired BC refractive index. Additional details on the sensitivity analysis can be found in Section 7. We have also expanded the discussion of the sensitivity of hybrid BNN predictions to assumed BC refractive index, coating refractive index, and monomer size in the discussion section (lines 474-482).

5. L130: Naming of input properties. For clarity, it may be helpful to explicitly distinguish the two groups of input features. For example, the first five inputs could be described as "microphysical/morphological descriptors," while Q_{ext}, SSA, and g could be referred to as "optical-property descriptors" or "physics-based optical descriptors." The authors do not necessarily need to follow these specific suggestions, but clearer categorization of the input features would improve readability. In particular, I found it somewhat difficult to follow the transition from optical-property-based inputs to the more general term "particle properties."

   We agree that describing the input features as "microphysical descriptors" and "homogeneous sphere optical descriptors" will help make the discussion more clear. We have added these labels to Figure 2 of the revised text and changed references to these variables where appropriate.

6. Sections 5 & 6: These results appear to demonstrate the superiority of the proposed hybrid BNN over the homogeneous-sphere and core-shell assumptions. However, this outcome seems somewhat expected because the DDA simulations are treated as the reference truth, and the hybrid BNN is directly trained on those DDA-generated optical properties. Therefore, Figures 3–5 may reflect the superiority of morphology-resolving DDA physics itself rather than the unique advantage of the hybrid BNN architecture. I wonder if the authors could clarify more explicitly what aspects of the improvement should be attributed to the Bayesian/surrogate framework versus the underlying DDA-based physical representation.

   We agree that the accuracy of hybrid BNN architecture is reflective of the accuracy of the underlying ADDA-generated dataset that was used to train the model, and therefore the superiority of BNN-predicted optical properties is rooted in the morphology-resolving ADDA physics. The advantage of the hybrid BNN model is that it is able to accurately predict the optical properties that would be calculated using ADDA without running the computationally-expensive ADDA code with associated uncertainty. We have explained this nuance in more detail in the revised manuscript on lines 265-273.

7. Figure 5: The SHAP analysis suggests that the homogeneous-sphere optical-property inputs contribute substantially to the hybrid BNN predictions. Since the framework conceptually combines homogeneous-sphere optical information with morphology-aware particle descriptors, it would be very helpful to more directly quantify the relative influence of these two groups of inputs on the final predictions. For example, an analysis comparing the hybrid BNN predictions with the corresponding homogeneous-sphere optical properties could help clarify the extent to which morphology-dependent particle information contributes beyond the baseline optical physics.

   We thank the reviewer for this thoughtful suggestion. Whereas Figure 5 identifies the influence of individual features, it does not directly quantify the relative contributions of the two feature groups. We, therefore, expanded the discussion in Section 6 to compare the relative importance of the combined microphysical descriptors with the homogeneous sphere optical descriptors on the target variables. We now report the aggregated mean absolute SHAP values for (1) the five microphysical descriptors and (2) the three homogeneous-sphere optical descriptors in lines 388-395 of the revised manuscript, which reflect the total importance of each group on the model predictions.

8. Figure 6: In panels B, E, and H, the hybrid BNN appears to exhibit relatively larger errors within portions of the parameter range represented in the training dataset. I may be misunderstanding the interpretation of this figure, but I initially expected the prediction errors to be minimized near the training-domain conditions. Additional clarification on how these perturbed-parameter experiments relate to the original training manifold versus the shaded "training range" would help improve interpretation of the figure.

   We would like to thank the reviewer for bringing our attention to Figure 6. The original version of Figure 6 incorrectly duplicated panels D, E, and F in panels G, H, and I. The revised manuscript has rectified this error and panels G, H, and I now correctly show errors in calculated g. Nevertheless, panel E still shows a local increase in the SSA error at the RI_{BC} value used in the training dataset. The perturbed cases were drawn from the held-out evaluation set, and the model is trained over a multidimensional parameter space. Each curve in Figure 6 is based on ten held-out particles for which one physical parameter was systematically perturbed. The shaded region indicates the RI_{BC} value used during training, but it does not indicate that the complete combination of particle size, morphology, coating amount, coating absorptivity, and wavelength lies near a densely sampled region of the full multidimensional training space. Model performance at a given RI_{BC} value, therefore, also depends on the specific combination of other particle properties represented by the selected particles. The local variation in panel E likely reflects the particular combinations of particle properties represented by this relatively small evaluation ensemble rather than a systematic reduction in performance at the training value. Further, we note that, although the SSA

error in Figure 6E is relatively high at the training value of RI_{BC}, its magnitude remains below 0.05 and is smaller than the corresponding errors from the homogeneous-sphere and core-shell approximations. Thus, although this local variation is worth noting, the BNN maintains relatively small SSA errors across the tested RI_{BC} values.

9. Uncertainty decomposition and Monte Carlo sampling: The uncertainty decomposition is an important component of the proposed hybrid BNN framework, but the manuscript provides relatively limited examples and interpretation of its behavior across different particle conditions. Are there general patterns or tendencies in the aleatoric versus epistemic uncertainty components with respect to particle size, morphology, coating amount, or wavelength? In addition, how sensitive are the uncertainty estimates to the Monte Carlo sampling sizes used during inference? Some discussion regarding convergence behavior and the choice of sampling size would help clarify the robustness of the uncertainty characterization.

   Figure 7 shows epistemic uncertainty as a function of N_{pp} and V/V_0 for wavelengths that are commonly used in optical instruments. We have added a similar plot to the appendix which instead shows aleatoric uncertainty (Figure K1). The aleatoric and epistemic uncertainties follow broadly similar regions of elevated uncertainty, but the magnitude of aleatoric uncertainty is generally smaller than the epistemic uncertainty.

   The reviewer is correct that the number of Monte Carlo samples affects the uncertainty estimates. We have added Figure D1 to the Appendix showing the convergence of aleatoric and epistemic uncertainty (σ_a and σ_e, respectively) for different values of the nested Monte Carlo sample sizes S and L. We quantify convergence by comparing the estimated σ_a and σ_e values at each combination of S and L with the corresponding converged estimates obtained using large sample sizes ($S = L = 1000$). At the default values used in the main text, $S = L = 50$, the deviations from the $S = L = 1000$ reference case are within 4% for all target parameters and both uncertainty components. During inference, the computational cost scales as $O(S \times L)$. Therefore, increasing $S = L$ from 50 to 1000 would increase computational cost by approximately a factor of 400, while changing the uncertainty estimates by less than 4%. The inference-throughput measurements reported in Appendix E further quantify the computational cost of the chosen default. We therefore retain $S = L = 50$ during inference as a cost-accuracy compromise. We have added this reasoning to lines 238-242 of the main text and pointed readers to Figure D1 and Table D1 in the appendix.

10. Section 8: Since epistemic uncertainty is interpreted as uncertainty associated with limited training coverage, it would be interesting to examine how the epistemic and aleatoric uncertainty components change when the size of the training dataset is artificially reduced (e.g., using half of the current training samples). Such an experiment may help demonstrate whether the epistemic uncertainty behaves consistently with the proposed interpretation, and it may be possible without requiring additional DDA simulations.

    We agree that directly testing the proposed relationship between training coverage and epistemic uncertainty is important. We have therefore added Appendix Figure J1, which shows the results of a data-restriction experiment designed to test whether epistemic uncertainty responds to a reduction in local training coverage. We retrained the hybrid BNN after removing approximately 12% of the data, corresponding to the region $70 \leq N_{pp} \leq 180$ and $8 \leq V/V_0 \leq 25$. Test particles within this masked region were retained for evaluation but were not used during retraining. For these particles, epistemic uncertainty increased by 14.4%, 28.6%, and 11.9% for Q_{ext}, SSA, and g, respectively, relative to the baseline model trained using the full dataset. For test particles outside the masked region, the corresponding increases were 3.9%, 5.0%, and −1.4%. Thus, the increase in epistemic uncertainty was larger inside the masked region than outside for all three target variables, supporting the interpretation that epistemic uncertainty responds to reduced training coverage. We also examined the corresponding changes in

aleatoric uncertainty, which increased by 6.9%, 27.9%, and 3.6% inside the masked region for Q_{ext}, SSA, and g, respectively, compared with increases of 0.8%, 8.7%, and −0.4% outside the region. Because the aleatoric and epistemic components are learned jointly through the same predictive model, changes in the training distribution can affect both components. We therefore interpret this experiment primarily as a test of whether the epistemic component responds to reduced training coverage. The localized increase in epistemic uncertainty is evident for all three target variables, although the magnitude of the response is target-dependent. We have added this discussion to Section 8 of the revised manuscript (lines 425-429) and pointed readers to Appendix Figure J1.

11. From the perspective of future radiative transfer and remote sensing applications, it would be interesting to know whether the framework could eventually be extended to predict more detailed angular scattering information beyond the asymmetry parameter, such as single-particle phase functions or phase-matrix-related quantities. Although the current study focuses on single-particle optical properties, such extensions could further increase the applicability of the framework to aerosol optics studies.

    The scattering phase matrix of particles was also saved during DDA calculations for this work. Therefore, future versions of the hybrid BNN could be extended to predict angular scattering, Legendre polynomials, or depolarization. Such extensions are planned for future work. We have noted in the discussion section that such extensions are possible for future model versions (lines 490-491).

Reviewer 2:

Assessment:

The paper describes a new hybrid Bayesian neural network model trained on discrete dipole approximation (DDA) simulations to better simulate the optical properties of particles containing black carbon. The authors compare the performance of the model with homogeneous and coated-sphere Mie theory and find that the Bayesian NN model performs better when simulating several optical parameters. One of the particularly interesting aspects of the study is the ability of the model to also determine uncertainties and therefore understand which input parameter space in DDA should be explored further.

The paper is well written and mostly clear, and the study is detailed and impactful. After the authors address the reviewers' comments, I think the paper should be published.

General Comments:

1. The paper uses DDA as the truth, which is very reasonable; however, it would be interesting if the authors would also venture into providing some recommendations on what experimental laboratory or field studies should focus on in terms of "tuning" parameters, observational variables, scenarios or conditions, to fill similar gaps to those the authors identify in DDA simulation space.

    Based on a review of laboratory studies of the optical properties of BC-containing particles, we would recommend that future laboratory experiments target smaller aggregates ($N_{pp} < 100$) that are thinly coated ($V/V_0 < 5$). Previous laboratory experiments have shown that particles in this N_{pp} range undergo moderate coating-induced restructuring when $1 < V/V_0 < 5$ (Corbin et al., 2023). Therefore, we expect particles in this region to have large morphological variability. This is supported by high epistemic and aleatoric uncertainty for particles in this region. Therefore, measurements of optical properties for these particles would further constrain model estimates and reduce epistemic uncertainties while better

characterizing the resolved variability represented by the aleatoric component. Measurements that jointly constrain aggregate size, coating amount, morphology, and wavelength-resolved optical properties would be particularly valuable. We have included this discussion on lines 445-452 of the revised manuscript.

2. It would have been interesting to also add comparisons with the RDG approximation, at least for uncoated BC particles.

   In our original manuscript, we limited comparisons of hybrid BNN predictions to Mie theory in the main text, given that homogeneous-sphere and core-shell models are currently implemented in large-scale models. We have added a comparison between the hybrid BNN and the RDG approximation to Appendix Figure I1. This comparison shows that the hybrid BNN is more accurate than RDG approximation for Q_{ext}, with median errors between 0.02 and 0.05 for the hybrid BNN and between -1.4 and 0.46 for RDG approximation. The hybrid BNN and RDG approximation yielded similar predictions for SSA, with median errors between -0.01 and 0.02 and between -0.04 and -0.01, respectively. This comparison shows that the magnitude of scattering and absorption is underestimated by the RDG approximation, but the ratio of scattering to absorption is well-captured. Finally, the hybrid BNN is also more accurate than the RDG approximation when calculating g for coated aggregates, with median errors between -0.02 and 0.01 for the hybrid BNN and between -0.20 and -0.15 for RDG. We have added this comparison to Appendix I and added a brief discussion of the results to lines 326-330 of the revised manuscript.

3. Perhaps a stronger emphasis should have been given to the absorption because, in the atmosphere, the other optical parameters (for example, scattering, SSA and g) are dominated by particles other than those containing BC, owing to their lower concentrations; while it is the absorption properties of BC that are more radiatively important (e.g., to determine heating rates) despite their lower concentration due to their stronger absorption cross sections with respect to most of the other atmospheric particles. For example, in Figure 3, I would also add a plot with Q_{abs} (instead of relegating it to Figure E1), even if that might be somewhat less sensitive to D_f.

   The reviewer is correct that total scattering is usually dominated by BC-free particles, while BC-containing particles have a strong influence on absorption. However, we emphasize Q_{ext}, SSA, and g because these quantities are commonly supplied to radiative transfer models, while Q_{abs} is typically calculated from Q_{ext} and SSA. We have expanded our discussion of the relative insensitivity of absorption to particle shape in the revised text (lines 316-325), which directs readers to Appendix H for more information.

4. Some additional references on previous work on these issues, especially in the realm of experiments and field studies, would be helpful.

   We have added citations to Cappa et al. (2012), Hu et al. (2021), Lack and Cappa (2010), Liu et al. (2017), Liu et al. (2020), and Shamjad et al. (2012) to the revised manuscript. These studies provide additional laboratory and field evidence for the influence of coating amount, mixing state, and particle morphology on the optical properties of BC-containing particles.

Specific Comments:

1. Line 30. Restructuring can also happen during cloud processing.

   We have specified that our reference to coating in this sentence could also apply to water during cloud processing.

2. Beginning of section 2. While Ensemble II "represents scenarios where liquid coatings induce restructuring via surface tension and capillary forces", what does Ensemble I represent? Is this more consistent

with processes such as cloud-induced compaction, or cases where the coating (other than water), initially present on the particle, eventually evaporates?

These particles generally represent particles with decoupled V/V_0 and $D_{f,c}$. The reviewer is correct that this could arise from coating evaporation. This could also arise if coating accumulates on BC aggregates via coagulation or via nanodroplet activation, which can occur under high supersaturation conditions. We have added this clarification to lines 74-77 of the revised manuscript.

3. Title of section 2.3: Although I do understand the reason for this wording choice, "numerically exact" might seem a contradiction with the term "approximation" in reference to the numerical model used (DDA). Perhaps a different section title would be better.

   We have changed the title of this section to "Optical Properties of Simulated Black Carbon-Containing Particles" in the revised manuscript.

4. Line 120: "other similar models" such as?

   We have specified in the revised manuscript that similar models include traditional neural networks which provide point-estimates of target variables.

5. The meaning of aleatoric vs. epistemic uncertainty is relatively clear; however, it would be even clearer if a simple example for each were provided.

   We have added a simple example which demonstrates aleatoric and epistemic uncertainty on lines 183-190 of the revised manuscript.

6. Line 159: Does theta include W and b (weights and biases)? If so, I find it confusing that later (in several following points in the paper) theta is referred to as " weights" only.

   We have changed references to θ as "weights" to "weights and biases" in the revised manuscript.

7. Lines 188-189: What is the reason for the miscalibration?

   The fitted values $\tau < 1$ indicate that the raw predictive intervals are systematically wider than required to achieve the observed validation set coverage, as described in Appendix C. These larger intervals may be associated with the independent-parameter approximation used for the Bayesian-layer parameters, which does not represent correlations between parameters, together with the nonlinear propagation from latent to physical space (Foong et al., 2020). We do not attribute the observed miscalibration to a single mechanism. Instead, we correct the observed miscalibration using post-hoc, per-target uncertainty scaling fitted on the validation set, following established approaches (Guo et al., 2017; Kuleshov et al., 2018). This scaling does not modify the trained model parameters and is applied only during inference. We clarify these possible sources of miscalibration and the subsequent calibration procedure on lines 174-175 and 207-208.

8. Line 192 and following: So, is tau just a tuning knob to match the validation targets?

   The reviewer's understanding is correct: τ is a post-hoc uncertainty-scaling factor, fitted on the held-out validation set after training by minimizing the Gaussian negative log-likelihood in latent space, following the process of Guo et al. (2017) and Kuleshov et al. (2018). The scaling is applied only at inference and does not modify the trained model parameters. Its purpose is to correct the small systematic miscalibration of raw variational-BNN predictive intervals described above. We have expanded the description of τ in the revised manuscript to highlight its post-hoc, calibration-only role (lines 207-217).

9. Line 208: How is the Gaussian noise chosen?

   We use Gaussian noise during inference as the model is trained using a diagonal Gaussian likelihood in latent space. The network predicts the mean and aleatoric standard deviation, and the standard normal noise is scaled by this predicted standard deviation during sampling. This makes the inference procedure consistent with the training objective. We have clarified the Gaussian sampling procedure and its relationship to the training likelihood in Section 3.3.2 of the revised manuscript.

10. Line 248: The distinction between D_f and $D_{f,c}$ is important. Can the authors explain how the two are calculated? Related to this point, on line 250, it is mentioned that $D_f \leq 2$ represents uncoated or coated particles not significantly compacted. Is the compaction here referring to the core alone or to the full particle? If it refers to the core only, as it would seem more physically appropriate, then the coating could still result in a particle with an overall D_f much higher than 2 even when the core is little or not at all compacted; for example, in the case where the coating is so thick that the information on the underlying core morphology is lost and the coating looks close to spherical on the outside surface. Some clarification on these points would be useful.

    The reviewer is correct that compaction here is in reference to the BC core. The reviewer is also correct that either a compacted core with a thin coating or an uncompacted core with a thick coating can result in $D_{\mathrm{f}} > 2$. We have added text from 276-282 of the revised text clarifying this point. We have also specified on lines 275-276 that D_{f} and $D_{\mathrm{f,c}}$ are calculated directly from the structure of each particle using the structure factor, which is described in detail in Sorensen (2001).

11. Line 255: This asymptotic behavior is reasonable and has been shown in the literature before.

    We clarified on lines 285-292 of the revised text that Mie theory has demonstrated accuracy for aged BC-containing particles with surface morphologies that approach that of a sphere. We also note that the hybrid BNN is also accurate for these particles, but has lower errors for particles with $D_{\mathrm{f}} < 2.5$.

12. Lines 282 – 283: As written, this sentence seems to imply more of a numerical or mathematical reason for the low sensitivity of Q_{abs} to D_f, just resulting from error cancellations in the calculations of Q_{ext} and SSA. However, I believe the low sensitivity of Q_{abs} to D_f has a much more physical foundation, at least for uncoated particles; a reason that is also at the root of why the simple RDG approximation (based on compounded monomer number or total mass) for Q_{abs} works relatively well as compared to Q_{sca} and therefore also for Q_{ext} and SSA. It is true that in the model here, Q_{abs} is given by Q_{ext} and 1-SSA to account for conservation of energy, but the physical lower sensitivity of Q_{abs} is reflected, not caused by this mathematical relation.

    We have modified lines 320-325 to reflect the physical basis for the accuracy of core-shell and homogeneous sphere approximations for calculating Q_{abs}. While Q_{abs} is often calculated from Q_{ext} and SSA, we agree that the appropriateness of equivalent-sphere approximations has a physical basis. This nuance is now reflected in the revised manuscript.

13. Line 298: This has indeed been shown to be the case in previous literature.

    We have noted that this effect has been demonstrated previously and referenced Sorensen's comprehensive 2001 review on the topic.

14. Line 299: Define m_{eff}.

    We have added a definition for m_{eff} on lines 345-349 of the revised manuscript.

Community Commenter 1:

Specific Comments:

1. The authors have performed a lot of DDA simulations and carefully described the training and validation starting from the DDA dataset (available at Zenodo). However, I couldn't find a description of the ADDA parameters used for simulations. It would be great to specify them for the overall reproducibility, including the ADDA version, DDA formulation, and discretization. Mentioning that some parameters are set to default values will also be fine.

   We have added Appendix A to the revised manuscript, which outlines the discretization and specifies the version of ADDA that was used.

2. Putting all raw DDA output online would probably be an overkill, but the authors may consider sharing the scripts, which perform these runs (to build a dataset). They will necessarily contain all the command line options for ADDA. Total or per run computational requirements of the DDA can also be interesting for readers.

   We have added an overview of timing requirements for our simulations in Appendix A and specified on lines 465-467 of the revised manuscript that the hybrid BNN speeds up computation time by 2-6 orders of magnitude. We have also added a tutorial covering input generation and ADDA scripting to the Zenodo repository.

3. Related question is that of the uncertainty of the DDA simulations (expected errors). Is there any estimate on that, at least for a few representative cases? It can be obtained, e.g., by several steps of discretization refinement for the same problem. I suspect, that the corresponding DDA errors are small enough to not influence any conclusions or further steps in the manuscript. Still, it would be great to quantify them. The latter seems especially natural, since the authors already perform an advanced error analysis, separating aleatoric and epistemic uncertainties. The DDA uncertainty probably fall into aleatoric class, but I am not sure if it needs any special consideration.

   The ADDA manual recommends maintaining a dipole length (d) small enough such that the number of dipoles per wavelength (λ/d) is greater than 10, with expected errors within a few percent. We use a dipole length of 0.0033 μm, which leads to a minimum λ/d value of 90 in our simulations. We therefore expect errors due to discretization to be very small.

   We also average the optical properties of each particle over 150 random orientations and train the hybrid BNN on the orientation-averaged optical properties. We compared the 150-orientation optical properties with 300-orientation optical properties of 40 randomly selected particles to quantify expected errors from orientation-averaging. We found that using 150 orientations in averaging led to absolute errors of 0.1% or less for all optical properties when compared to average values for 300 orientations (i.e. 0.1% improvement for double computational cost). We therefore expect negligible errors to be introduced due to discretization and orientation-averaging in ADDA simulations. We have added Appendix A to the revised manuscript discussing these details.

4. Finally, I have a minor stylistic note concerning the name of the ADDA code. The official guideline is not to deabbreviate it - see https://github.com/adda-team/adda/wiki/FAQ#what-is-the-official-name-of-the-code-what-does-a-stands-for. In other words, the standard naming is just ADDA. Similar aspects were recently discussed in another GMD paper - https://doi.org/10.5194/gmd-19-887-2026.

We have removed all instances of de-abbreviation in the revised text and ensured that all references to simulations use the proper notation "ADDA".

References

Beeler, P. and Chakrabarty, R. K.: Constraining the particle-scale diversity of black carbon light absorption using a unified framework, Atmospheric Chemistry and Physics, 22, 14 825–14 836, 2022.

Beeler, P., Corbin, J. C., Sipkens, T. A., and Fierce, L.: A Framework for Quantifying the Size and Fractal Dimension of Compacting Soot Particles, Environmental science & technology, 59, 5994–6003, 2025.

Cappa, C. D., Onasch, T. B., Massoli, P., Worsnop, D. R., Bates, T. S., Cross, E. S., Davidovits, P., Hakala, J., Hayden, K. L., and Jobson, B. T.: Radiative absorption enhancements due to the mixing state of atmospheric black carbon, Science, 337, 1078–1081, 2012.

Corbin, J. C., Modini, R. L., and Gysel-Beer, M.: Mechanisms of soot-aggregate restructuring and compaction, Aerosol Science and Technology, 57, 89–111, 2023.

Foong, A. Y. K., Burt, D. R., Li, Y., and Turner, R. E.: On the Expressiveness of Approximate Inference in Bayesian Neural Networks, https://arxiv.org/abs/1909.00719, 2020.

Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q.: On calibration of modern neural networks, in: International conference on machine learning, pp. 1321–1330, PMLR, 2017.

Hu, K., Liu, D., Tian, P., Wu, Y., Deng, Z., Wu, Y., Zhao, D., Li, R., Sheng, J., Huang, M., et al.: Measurements of the diversity of shape and mixing state for ambient black carbon particles, Geophysical Research Letters, 48, e2021GL094 522, 2021.

Kuleshov, V., Fenner, N., and Ermon, S.: Accurate uncertainties for deep learning using calibrated regression, in: International conference on machine learning, pp. 2796–2804, PMLR, 2018.

Lack, D. and Cappa, C.: Impact of brown and clear carbon on light absorption enhancement, single scatter albedo and absorption wavelength dependence of black carbon, Atmospheric Chemistry and Physics, 10, 4207–4220, 2010.

Liu, D., Whitehead, J., Alfarra, M. R., Reyes-Villegas, E., Spracklen, D. V., Reddington, C. L., Kong, S., Williams, P. I., Ting, Y.-C., and Haslett, S.: Black-carbon absorption enhancement in the atmosphere determined by particle mixing state, Nature Geoscience, 10, 184–188, 2017.

Liu, F., Yon, J., Fuentes, A., Lobo, P., Smallwood, G. J., and Corbin, J. C.: Review of recent literature on the light absorption properties of black carbon: Refractive index, mass absorption cross section, and absorption function, Aerosol Science and Technology, 54, 33–51, 2020.

Shamjad, P., Tripathi, S., Aggarwal, S., Mishra, S., Joshi, M., Khan, A., Sapra, B., and Ram, K.: Comparison of experimental and modeled absorption enhancement by black carbon (BC) cored polydisperse aerosols under hygroscopic conditions, Environmental science & technology, 46, 8082–8089, 2012.

Sorensen, C.: Light scattering by fractal aggregates: a review, Aerosol Science & Technology, 35, 648–687, 2001.