DDA-BNN v1.0: A Morphology-Aware Surrogate Model for the Optical Properties of Black Carbon–Containing Particles

Payton Beeler[1], Sam Donald[2], and Laura Fierce[1]

[1]Atmospheric, Climate, and Earth Sciences Division, Pacific Northwest National Laboratory, Richland, WA, 99354, USA
[2]Physical and Computational Sciences, Pacific Northwest National Laboratory, Richland, WA, 99354, USA

Correspondence: Laura Fierce (laura.fierce@pnnl.gov)

Abstract. Black carbon (BC) is the most strongly absorbing component of atmospheric aerosol and significantly impacts Earth's energy balance. The optical properties of BC-containing particles depend on particle-level variability in size, chemical composition, and internal morphology. Such particle-level details are not easily represented in large-scale atmospheric models. Existing parameterizations typically assume idealized particle geometries (e.g., homogeneous spheres or concentric core–shell spheres) and homogeneous mixing, which can yield biased predictions and provide no quantitative estimate of model-form uncertainty at the single-particle level. In this work, we present a probabilistic framework for predicting the optical properties of individual BC-containing particles using a hybrid Bayesian neural network (BNN) model trained on discrete dipole approximation (DDA) simulations. The hybrid BNN is a flexible combination of deterministic and Bayesian layers allowing for more realistic treatment of particle optical properties and quantification of uncertainty. The hybrid BNN model predicts extinction efficiency, single-scattering albedo, and asymmetry parameter and returns predictive uncertainty that can be decomposed into aleatoric (data-driven variability) and epistemic (uncertainty due to limited training coverage) components. We show that the hybrid BNN outperforms homogeneous-sphere and core–shell Mie approximations for calculating extinction and scattering-sensitive quantities (Q_{ext}, SSA and g), while maintaining comparable accuracy for absorption-related metrics. We further demonstrate how epistemic uncertainty highlights under-sampled regions of particle parameter space, enabling targeted design of future DDA simulations that most effectively reduce model uncertainty. This uncertainty-aware surrogate provides a practical pathway for incorporating realistically complex particle morphologies into parameterizations of aerosol optical properties, which will ultimately improve the reliability of model-based assessments of BC impacts on the atmosphere.

1 Introduction

Predicting the optical properties of black carbon (BC)-containing particles presents a unique and longstanding challenge in atmospheric models and remote sensing algorithms (Bond and Bergstrom, 2006). BC-containing particles represent a small fraction of atmospheric aerosol mass but contribute significantly to total light absorption (Bond et al., 2013), so modest errors in BC optics can translate into disproportionately large errors in aerosol direct radiative effect and inferred emissions. Despite decades of study, a central problem remains: particle-resolved morphology and mixing state strongly control absorption and

scattering, yet large-scale models require computationally cheap optics schemes that reduce this complexity into a small set of bulk parameters.

BC-containing particles are distinct from many other aerosol types because of their strong and highly variable absorption properties (Samset et al., 2018; Laskin et al., 2015; Bond and Bergstrom, 2006; Cappa et al., 2012; Fierce et al., 2020; Liu et al., 2017b, 2
. BC-containing particles are emitted as fractal aggregates of primary spherules, and their morphology evolves after emission through coagulation and condensation of semi-volatile materials (China et al., 2013). Accumulation of coatings (or water during cloud processing) can restructure aggregates toward more compact configurations via surface tension and capillary forces (Corbin et al., 2023; Beeler et al., 2025), while the coating itself can range from weakly absorbing to strongly absorbing ("brown carbon", BrC). Laboratory and ambient observations have shown these combined changes in internal structure, coating amount, and coating absorptivity lead to broad variability in single-particle extinction, single-scattering albedo (SSA), and asymmetry parameter (g) (Hu et al., 2021; Liu et al., 2020; Bond et al., 2013).

To achieve computational efficiency, large-scale atmospheric and radiative transfer models typically represent BC-containing particles using simplified shapes (homogeneous spheres or core–shell spheres) and assume uniform composition within size modes (Liu et al., 2016). While these assumptions simplify calculations, they introduce substantial error by neglecting complex particle structures and compositional diversity (Fierce et al., 2016, 2020; Liu et al., 2015). Detailed models of single-particle light scattering, such as those based on the discrete dipole approximation (DDA), represent realistically complex particle morphologies and compositions (Yurkin and Hoekstra, 2011), but are too computationally expensive for routine use in regional or global modeling and are impractical for embedding directly into iterative workflows such as inverse modeling, data assimilation, or uncertainty propagation. Recent work has explored the use of machine learning (ML) to approximate optics calculations for BC-containing particles with much lower computational cost (Qin et al., 2025; Romshoo et al., 2024; Li et al., 2023; Li and May, 2022). However, these studies rely on multiple-sphere T-matrix models, which are more computationally efficient than DDA, but generally require idealized representations (e.g., non-overlap constraints and idealized coatings) (Liu et al., 2017a). Additionally, most existing ML surrogates return only point predictions and do not separate uncertainty due to irreducible variability in the training data from uncertainty due to sparse coverage of particle parameter space. This is a key limitation because the relevant BC parameter space is high-dimensional and observational constraints remain incomplete.

As a path forward, this work introduces a hybrid Bayesian Neural Network (BNN) approach (Bonnet et al., 2023; Olivier et al., 2021; Gal and Ghahramani, 2016) to predict the absorption and scattering properties of individual BC-containing particles, as well as confidence intervals in modeled quantities. The hybrid BNN combines deterministic and Bayesian layers to balance computational efficiency with uncertainty quantification. The model is trained on single-particle optics simulations that have been informed by detailed laboratory measurements of the mixing state-dependent morphology of BC-containing particles (Beeler et al., 2025). The hybrid BNN predicts the extinction efficiency, single-scattering albedo (SSA), and asymmetry parameter (g) and associated uncertainties for individual particles, allowing for the decomposition of total model uncertainty into aleatoric and epistemic components. We compare the performance of the hybrid BNN against homogeneous-sphere and core–shell Mie approximations commonly used in atmospheric models. Finally, we demonstrate how epistemic uncertainty maps can be used to design targeted future DDA experiments that reduce surrogate uncertainty most efficiently. This hybrid

BNN lays the groundwork for an uncertainty-aware, computationally efficient surrogate to enable more complete representation of BC optical properties in larger-scale models.

2 Training Set Generation

The training dataset comprises two ensembles of simulated BC-containing particles designed to span contrasting representations of the coupling between particle coatings and aggregate morphology. Ensemble I ("non-compacting") is taken from Beeler and Chakrabarty (2022) and intentionally decouples coating amount from aggregate compaction, representing scenarios where coating does not induce restructuring. Ensemble II ("compacting") includes particles whose compaction depends on coating amount following laboratory-informed scaling relationships (Beeler et al., 2025). These particles represent scenarios where liquid coatings induce restructuring via surface tension and capillary forces (Corbin et al., 2023). Together, the two ensembles allow the surrogate to learn both morphology-driven and coating-driven variability while retaining physical plausibility for compacting particles. The following sections provide details of the particles in each ensemble and calculation of their optical properties.

2.1 Ensemble I: Non-Compacting Black Carbon-Containing Particles

Non-compacting BC-containing particles are taken from Beeler and Chakrabarty (2022), and full details on their generation can be found therein. Briefly, these particles have three morphologies (fresh, partially collapsed, fully collapsed), with each morphology ranging from uncoated to thickly coated (Figure 1). These particles are representative of BC-containing particles with decoupled morphologies and coating thicknesses. Physically, decoupled morphologies and coating thicknesses can arise if coatings evaporate from BC aggregates, coating is added via coagulation, or coating accumulates under high supersaturation conditions (Corbin et al., 2023). External coatings are quantified using the ratio of total particle volume to BC volume, also known as the volume growth ratio (V/V_0). The amount of BC in each particle is quantified by the number of primary particles in the aggregate (N_{pp}). All particles in this group are made up of primary particles with radii of 20 nm. The degree of compaction is quantified using the core fractal dimension ($D_{\mathrm{f,c}}$), with fresh aggregates having low $D_{\mathrm{f,c}}$ (1.8) and compact aggregates having high $D_{\mathrm{f,c}}$ (> 2.5) (Koeylue et al., 1995). The $D_{\mathrm{f,c}}$ of each aggregate is calculated from the structure factor. The structure factor, which governs the angular dependence of scattered light, is the Fourier transform of the density autocorrelation function. The structure factor exhibits an inverse power-law dependence as a function of the scattering vector in the fractal regime, with an exponent equal $D_{\mathrm{f,c}}$ (Sorensen, 2001).

It is important to note that $D_{\mathrm{f,c}}$ and V/V_0 are statistically independent in Ensemble I. We simulated 900 particles in Ensemble I, with N_{pp} ranging from $13-1189$, $D_{\mathrm{f,c}}$ ranging from $1.8-3.0$, and V/V_0 ranging from $1-75$ (see Table 1). The optical properties are calculated at wavelengths (λ) of 405, 532, 880, 1200, and 2500 nm. The refractive index of BC is assumed constant at $1.95+0.79i$, and the coating refractive index is $1.55+[0.00, 0.01, 0.05]i$ (Bond and Bergstrom, 2006).

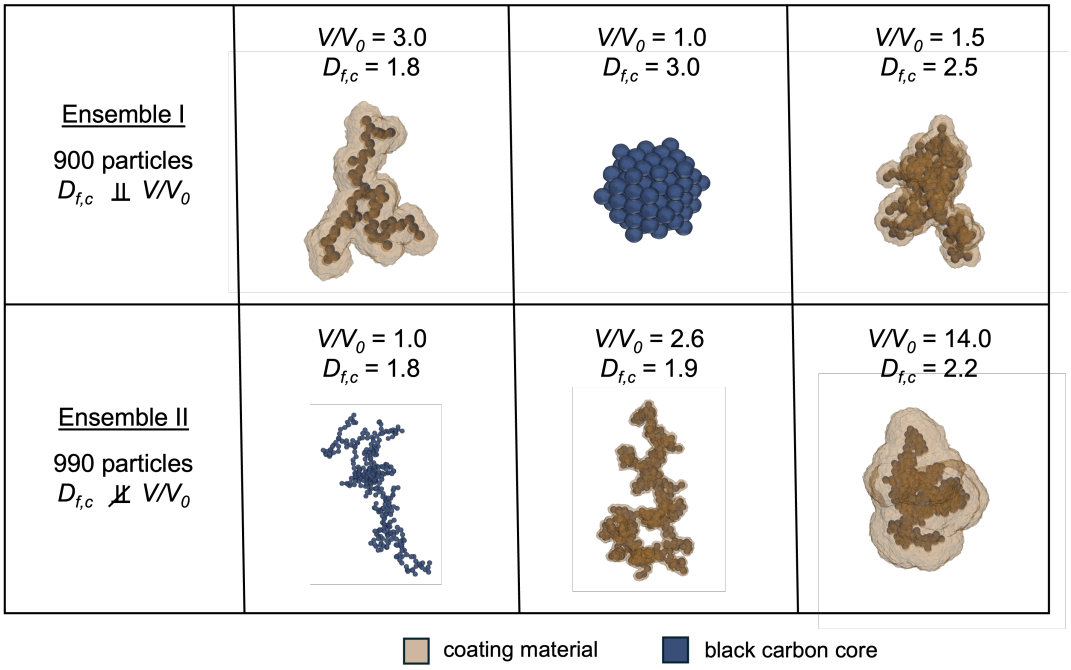


Figure 1. Examples of simulated particles used to train the hybrid Bayesian neural network model. The training dataset is made up of two ensembles of particles (I and II). Particles in Ensemble I (top row) are taken from a previous publication (Beeler and Chakrabarty, 2022) and have statistically independent (⫫) total-to-black carbon volume ratio (V/V_0) and core fractal dimension ($D_{\mathrm{f,c}}$). The particles in Ensemble II (bottom row) represent compacting particles and have V/V_0 that is calculated based on $D_{\mathrm{f,c}}$ following Beeler et al. (2025). Therefore, V/V_0 and $D_{\mathrm{f,c}}$ are not statistically independent for these particles ($D_{\mathrm{f,c}}$ ⫫̸ V/V_0).

2.2 Ensemble II: Compacting Black Carbon-Containing Particles

Previous studies have developed scaling laws for $D_{\mathrm{f,c}}$ as a function of V/V_0 using detailed laboratory measurements and single-particle simulations (Beeler et al., 2025). Therefore, the $D_{\mathrm{f,c}}$ of particles in this ensemble is dependent on V/V_0 (following the formulation of Beeler et al. (2025)). The N_{pp} of compacting particles is evenly sampled in log-space from 20 to 1000, and the V/V_0 of each particle is randomly selected between 1 and 100 (see Table 1). The $D_{\mathrm{f,c}}$ of a particle for a given V/V_0 is given by

$$D_{\mathrm{f,c}} - D_{f,0} = c_1\{\exp[c_2\,(x_c - x_s)] - \exp[c_2\,(x_0 - x_s)]\}, \tag{1}$$

where x_c is given by

$$\frac{x_c - x_f}{x_0 - x_f} = \exp\left[-0.799\left(\frac{V}{V_0} - 1\right)\right], \tag{2}$$

and x_f is given by

$$x_f = 1.07 x_0 N_{\mathrm{pp}}^{-0.026}. \tag{3}$$

Table 1. Overview of discrete dipole approximation (DDA) simulations used to train hybrid Bayesian neural network. The training dataset is made up of two ensembles. Particles in Ensemble I are taken from a previous publication and have independent total-to-black carbon volume ratio (V/V_0) and core fractal dimension ($D_{\mathrm{f,c}}$), while these are coupled for particles in Ensemble II. Other parameters used as inputs to the DDA model include the number of primary particles in each aggregate (N_{pp}), the radius of primary particles (r_{pp}), the wavelength of incident light (λ), the refractive index of BC ($\mathrm{RI_{BC}}$), the coating refractive index at 550 nm ($\mathrm{RI_{coat,550\ nm}}$), and the absorption Ångstrom exponent of the coating ($\mathrm{AAE_{coat}}$).

Ensemble	Data Points	N_{pp}	r_{pp} (nm)	V/V_0	$D_{\mathrm{f,c}}$	λ (nm)	$\mathrm{RI_{BC}}$	$\mathrm{RI_{coat,550\ nm}}$	$\mathrm{AAE_{coat}}$
I	900	$13-1189$	20	$1-75$	$1.8-3.0$	405, 532, 880, 1200, 2500	$1.95+0.79i$	$1.55+[0.00, 0.01, 0.05]i$	–
II	990	$20-1000$	20	$1-100$	$1.8-2.5$	$300-1200$	$1.95+0.79i$	$1.6+[0.00-0.25]i$	$0.35-7.6$

Here, $D_{f,0}=1.8$, $x_0=0.46$, $x_s=1/3$, $c_1=11.66$, and $c_2=-41.39$.

BC-containing particles with given N_{pp} are generated using diffusion-limited cluster-cluster aggregation, compaction to the desired $D_{\mathrm{f,c}}$ (Eq. 1) is simulated using a spring-van der Waals compaction model, and coating to the desired V/V_0 is simulated using a capillary coating model; each of which are described in detail in Beeler et al. (2025). The $D_{\mathrm{f,c}}$ of each particle in Ensemble II is then calculated using its structure factor.

Ensemble II included 990 particles that are representative of atmospheric BC-containing particles that have undergone coating-induced compaction. Half of the particles in Ensemble II are coated with non-absorbing coatings and have a coating refractive index of $1.6+0.0i$. The remaining half of Ensemble II particles are coated with brown carbon (BrC). The imaginary part of the BrC refractive index (κ) at the wavelength of interest is given by

$$\kappa(\lambda)=\kappa_{550\ \mathrm{nm}}\left(\frac{\lambda}{550\ \mathrm{nm}}\right)^{-\mathrm{AAE_{coat}}}. \tag{4}$$

Here, $\kappa_{550\ \mathrm{nm}}$ is the imaginary part of the refractive index at 550 nm and $\mathrm{AAE_{coat}}$ is the coating absorption Ångstrom exponent. Previous studies have shown that there is a relationship between $\kappa_{550\ \mathrm{nm}}$ and $\mathrm{AAE_{coat}}$, given by

$$\mathrm{AAE_{coat}}=7.6\exp\left[-12.4\kappa_{550\ \mathrm{nm}}\right]. \tag{5}$$

This relationship represents the measured properties of BrC across absorptivity (Saleh et al., 2018). The optical properties of Ensemble II particles are calculated at wavelengths randomly selected in 50 nm-spaced bins between 300 and 1200 nm, the refractive index of BC is assumed constant at $1.95+0.79i$, the radius of primary particles is 20 nm, $\kappa_{550\ nm}$ is randomly sampled between 0.0 and 0.25, and the coating refractive index at the chosen wavelength is given by equations 4 and 5.

2.3 Optical Properties of Simulated Black Carbon-Containing Particles

We use the ADDA code to calculate the optical properties of each particle in the training set (Yurkin and Hoekstra, 2011). This model operates by breaking each aggregate into sub-volumes that are much smaller than the wavelength of incident light (dipoles). Each dipole is then treated as a point scatterer that interacts with surrounding dipoles. This method has been shown to accurately calculate the optical properties of arbitrarily shaped objects and is highly accurate in capturing near-field interactions among the dipoles present in touching BC monomers. A full description of the ADDA configuration can be found in Appendix A. The optical properties of particles in Ensemble I are calculated at 405, 532, 880, 1200, or 2500 nm. We calculate the optical properties of Ensemble II particles at wavelengths randomly distributed among 50 nm-spaced bins ranging from 300 to 1200 nm.

3 Hybrid Bayesian Neural Network Description

Accurate prediction of the optical properties of BC-containing particles is of the utmost importance for determining their behavior in and effects on the atmosphere. However, BC-containing particles have variable compositions and morphologies, both of which have a large but uncertain effect on their optical properties. Here, we train a hybrid BNN model to infer the optical properties of BC-containing particles coated by absorbing and non-absorbing materials, at different stages of compaction. This model is unique from other similar models (such as traditional neural networks) as it enables uncertainty quantification of the inferred optical properties. Uncertainty quantified by the hybrid BNN can further be separated into aleatoric (irreducible noise inherent to the training dataset) and epistemic (error due to limited data). Uncertainty quantification in single-particle optical properties can then be propagated to particle populations, allowing for more complete representation of BC optical properties in atmospheric models. The hybrid BNN framework developed here is a sequence of deterministic and Bayesian layers trained on the dataset detailed in Section 2. The following sections outline data pre-processing routines, model architecture, uncertainty decomposition, and model evaluation.

3.1 Data Pre-Processing

Because several optical properties of BC-containing particles are correlated with one another (such as absorption and extinction efficiencies), we have chosen the extinction efficiency (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) as the target variables of the hybrid BNN model. These optical properties are also commonly supplied to radiative transfer algorithms. Similarly, several properties of BC-containing particles are also correlated with one another, such as volume-equivalent diameter and particle mass. To avoid potential correlations, we have selected five microphysical properties of BC-containing particles as input variables to the hybrid BNN model: N_{pp}, V/V_0, $D_{\mathrm{f,c}}$, the volume-equivalent size parameter (χ_{ve}), and the imaginary part of the coating refractive index (κ_{coat}). These inputs are chosen as they are either directly calculated or easily derived by large-scale atmospheric and radiative transfer models. Additionally, the collection of these parameters fully describes the amount of BC (N_{pp}), the amount of coating (V/V_0), the shape of particles ($D_{\mathrm{f,c}}$), the coating material (κ_{coat}),

and the wavelength of light (χ_{ve}). The optical properties of a volume-equivalent homogeneous sphere are also used as input variables in the hybrid BNN ($Q_{\text{ext,hs}}$, SSA_{hs}, and g_{hs}). The optical properties of a homogeneous volume-equivalent sphere are calculated using Mie theory, with refractive index given by the volume-weighted average of the BC and coating refractive indices (Sumlin et al., 2018). These were used as input variables because they are strong predictors of the target variables and are currently calculated in large-scale models.

Prior to training, each component of the target vector $y_t = (Q_{\text{ext}}, SSA, g)$ is mapped to an unbounded real domain using a transformation T that preserves physical bounds (i.e., $T(y_t) = (T(Q_{\text{ext}}), T(SSA), T(g))$). For strictly positive targets (e.g., optical efficiencies), we use a logarithmic transform. For bounded targets in $(0,1)$ (SSA and g) (Koshelev, 2015), we use the logit transform. While the theoretical limits of g are $(-1,1)$, it is extremely rare for atmospheric particles to have $g < 0$ (Zhao et al., 2018). Therefore, we use the logit transform for g. Specifically, T is given by:

$$T(y) = \begin{cases} \ln(y), & y \in (0,\infty) \\ \text{logit}(y) = \ln\left(\frac{y}{1-y}\right), & y \in (0,1) \end{cases} \tag{6}$$

These transformations preserve constraints of the underlying physical variables while allowing the model to operate in an unconstrained latent space to ensure appropriately bounded uncertainties. After T transformations have been applied, all input variables (x) defined above and transformed target variables (y_t) are independently standardized using Z-score normalization (Fei et al., 2021) to give inputs and target variables in unconstrained latent space ($\tilde{x}$ and $\tilde{y_t}$, respectively). More details on Z-score normalization can be found in Appendix B.

3.2 Model Architecture

The proposed neural network (NN_θ) is made up of an arbitrary sequence of deterministic (D) or Bayesian (B) layers (Kendall and Gal, 2017; Gal and Ghahramani, 2016). For example:

$$NN_\theta = [D, B, D], \tag{7}$$

defines a three-layer network with a deterministic input layer, a Bayesian hidden layer, and a deterministic output layer. Each layer (z) is generally defined as

$$z \mapsto \phi(Wz + b), \tag{8}$$

where W and b are the model weights and biases, and ϕ is the activation function. In a deterministic layer, $\theta \in (W, b)$ are the fixed parameters learned during training. Conversely, in a Bayesian layer the weights and biases are treated as random variables with an approximate (variational) posterior distribution, e.g., $q_\phi(\theta) = \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, parameterized by a learnable mean μ_θ and scale (standard deviation) σ_θ, from which θ is sampled during training and inference. We model each Bayesian-layer parameter independently, so this approximation does not represent correlations between parameters.

The outputs of the model, for a given set of sampled weights and biases θ and latent-space inputs $\tilde{x}$, are a predicted mean vector $\tilde{\mu}_\theta(\tilde{x})$ and a vector of per-target aleatoric standard deviations $\tilde{\sigma}_{a,\theta}(\tilde{x})$ in latent space. These outputs are mapped back

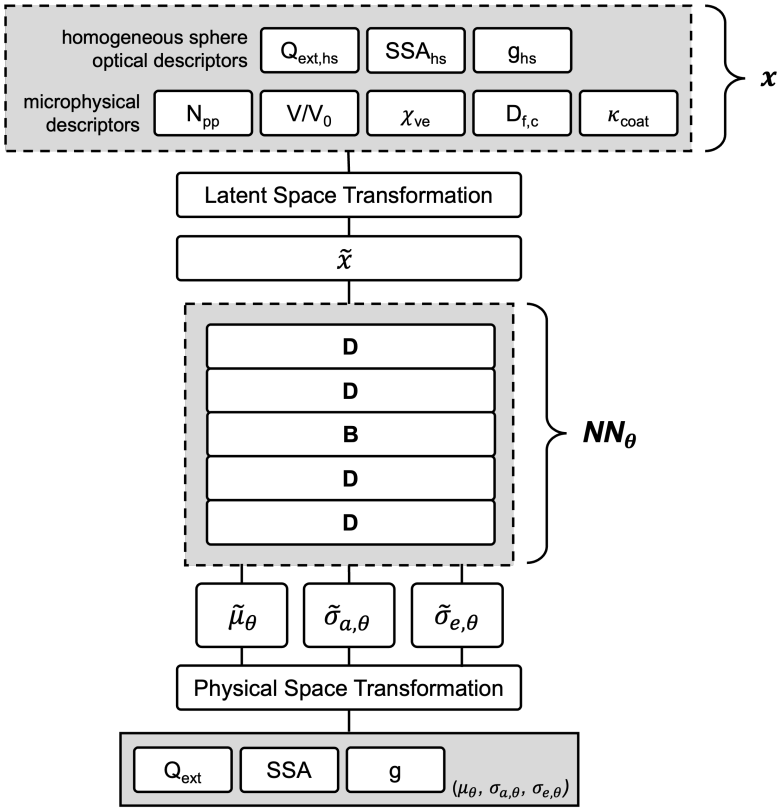


Figure 2. Example of the hybrid Bayesian neural network architecture. Inputs are standardized using Z-score normalization, and target variables are mapped to an unbounded latent space using the transform T followed by Z-score normalization. The neural network (NN_θ) consists of a combination of deterministic and Bayesian layers producing a latent-space mean ($\tilde{\mu}_\theta$), aleatoric uncertainty ($\tilde{\sigma}_{a,\theta}$), and epistemic uncertainty ($\tilde{\sigma}_{e,\theta}$). Predictions are propagated to physical space by sampling in latent space and then applying the inverse transforms (Z^{-1} and T^{-1}) to obtain optical properties in physical units.

to physical units via the inverse Z and T transforms (Section 3.1) using Monte Carlo sampling (Section 3.3). We assume a diagonal aleatoric covariance (independent targets) and leave prediction of the full aleatoric covariance matrix for future work. The proposed pipeline is visualized in detail in Figure 2.

3.3 Uncertainty Decomposition

During inference we seek both a point prediction and a decomposition of predictive uncertainty into aleatoric $\tilde{\sigma}_a$ and epistemic $\tilde{\sigma}_e$ components. Suppose two observers are trying to predict the outcome of the next flip of an unknown coin. The first observer has recorded the results of 1,000 previous flips and estimates that the probability of heads is 0.6. Their remaining uncertainty about the next flip is dominated by aleatoric uncertainty; even with perfect knowledge of the coin's bias, each flip is still random. The second observer has not seen any previous flips and does not know whether the coin is fair or biased. Their

uncertainty includes epistemic uncertainty arising from their lack of knowledge about the coin, in addition to the aleatoric uncertainty inherent in the flipping process. Therefore, aleatoric uncertainty represents irreducible variability in the data (e.g., unresolved physical variability and numerical noise in the training set), whereas epistemic uncertainty reflects uncertainty in the model parameters due to limited data and/or model capacity. The initial aleatoric and epistemic decompositions are computed in latent space under Gaussian assumptions (Section 3.3.1). We then calibrate the uncertainty (Section 3.3.2) before propagating to physical space by sampling and applying the inverse transforms, from which we report predictive quantiles in physical units and estimate physical-space aleatoric and epistemic standard deviations via a nested Monte Carlo approximation (Section 3.3.3).

3.3.1 Latent-Space Predictive Distribution

At inference time we draw S independent weight and bias samples $\theta^{(s)} \sim q_\phi(\theta)$ from the learned approximate posterior. For each draw, the neural network NN_θ directly outputs $\tilde{\mu}^{(s)}(\tilde{x}) := \tilde{\mu}_{\theta^{(s)}}(\tilde{x})$ and $\tilde{\sigma}_a^{(s)}(\tilde{x}) := \tilde{\sigma}_{a,\theta^{(s)}}(\tilde{x})$. The latent-space predictive mean is estimated by Monte Carlo averaging,

$$\tilde{\mu}_\theta(\tilde{x}) = \frac{1}{S}\sum_{s=1}^{S} \tilde{\mu}^{(s)}(\tilde{x}). \tag{9}$$

Aleatoric uncertainty in latent space is estimated as the mean of the predicted variances,

$$\tilde{\sigma}^2_{a,\theta}(\tilde{x}) = \frac{1}{S}\sum_{s=1}^{S} \left[\tilde{\sigma}_a^{(s)}(\tilde{x})\right]^{\odot 2}, \tag{10}$$

and epistemic uncertainty is estimated from the variations of the sampled means across draws,

$$\tilde{\sigma}^2_{e,\theta}(\tilde{x}) = \frac{1}{S}\sum_{s=1}^{S} \left[\tilde{\mu}^{(s)}(\tilde{x}) - \tilde{\mu}_\theta(\tilde{x})\right]^{\odot 2}. \tag{11}$$

Finally, the total latent-space predictive uncertainty is taken as the sum of epistemic and aleatoric variances,

$$\tilde{\sigma}^2_{tot,\theta}(\tilde{x}) = \tilde{\sigma}^2_{a,\theta}(\tilde{x}) + \tilde{\sigma}^2_{e,\theta}(\tilde{x}). \tag{12}$$

3.3.2 Uncertainty Scaling

Although Eqs. (10–11) provide an aleatoric-epistemic decomposition in latent space, the overall predictive uncertainty can be miscalibrated due to non-linear activation functions, leading to prediction intervals that are systematically too wide or narrow. We therefore apply per-target uncertainty scaling in latent space, using multiplicative factors τ fitted on a validation set (Guo et al., 2017; Kuleshov et al., 2018).

We fit τ by minimizing the Gaussian negative log-likelihood in latent space, using the latent predictive mean $(\tilde{\mu}_\theta(\tilde{x}))$, the total latent predictive standard deviation $(\tilde{\sigma}_{\mathrm{tot},\theta}(\tilde{x}))$, and the validation targets transformed into latent space $(\tilde{y}_\mathrm{t} = ZT(y_\mathrm{t}))$. For each target dimension j, we solve

$$\tau_j = \arg\min_{\tau>0} \frac{1}{N}\sum_{i=1}^{N} \mathrm{NLL}(\tilde{y}_{\mathrm{t},j}(x_i)\,;\, \tilde{\mu}_{\theta,j}(\tilde{x}_i), \tau\,\tilde{\sigma}_{\mathrm{tot},\theta,j}(\tilde{x}_i)). \tag{13}$$

At inference time, the calculated τ values are applied in latent space before propagation to physical space, scaling both epistemic and aleatoric uncertainty components. Details of the fitted scaling factors along with their effect on predictive-interval coverage and reliability are reported in Appendix C.

3.3.3 Propagation to Physical Space

Because T^{-1} is nonlinear, the assumed Gaussian distribution in latent space does not necessarily map to a Gaussian distribution in physical space. We therefore propagate uncertainty by Monte Carlo sampling in latent space and applying inverse transforms. Specifically, starting from the S posterior draws $\{\tilde{\mu}^{(s)}(\tilde{x}), \tilde{\sigma}_a^{(s)}(\tilde{x})\}_{s=1}^{S}$, we apply the latent-space scaling factors τ and then generate physical-space predictive samples via nested Monte Carlo sampling. To scale the epistemic component (the spread across S), we keep the Monte Carlo mean fixed and scale the deviations around it,

$$\tilde{\mu}_\tau^{(s)}(\tilde{x}) = \tilde{\mu}_\theta(\tilde{x}) + \tau \odot \Big(\tilde{\mu}^{(s)}(\tilde{x}) - \tilde{\mu}_\theta(\tilde{x})\Big). \tag{14}$$

We then directly scale the corresponding per-sample aleatoric standard deviations.

$$\tilde{\sigma}_{a,\tau}^{(s)}(\tilde{x}) = \tau \odot \tilde{\sigma}_a^{(s)}(\tilde{x}). \tag{15}$$

Next, for each posterior draw s we generate L latent samples by adding Gaussian noise,

$$\tilde{y}^{(s,\ell)}(\tilde{x}) = \tilde{\mu}_\tau^{(s)}(\tilde{x}) + \tilde{\sigma}_{a,\tau}^{(s)}(\tilde{x}) \odot \eta^{(s,\ell)}, \qquad \eta^{(s,\ell)} \sim \mathcal{N}(0, I), \quad \ell = 1, \ldots, L. \tag{16}$$

Finally, each latent sample is mapped back to physical space,

$$y^{(s,\ell)}(x) \;=\; T^{-1}\Big(Z^{-1}\big(\tilde{y}^{(s,\ell)}(\tilde{x})\big)\Big), \tag{17}$$

with predictive summaries (e.g., means, percentiles) computed from the resulting physical-space samples $\{y^{(s,\ell)}(x)\}$.

We then estimate aleatoric and epistemic uncertainties in physical units using a nested Monte Carlo approximation to the law of total variance (Kendall and Gal, 2017). Let $\mathbb{E}_\ell[\cdot]$ and $\mathrm{Var}_\ell[\cdot]$ denote the mean and variance over $\ell = 1, \ldots, L$ at fixed s, and let $\mathbb{E}_s[\cdot]$ and $\mathrm{Var}_s[\cdot]$ denote the mean and variance over $s = 1, \ldots, S$. We compute

$$\widehat{\mathrm{Var}}_a(y \mid x) \approx \mathbb{E}_s\Big[\mathrm{Var}_\ell\Big(y^{(s,\ell)}(x)\Big)\Big], \tag{18}$$

$$\widehat{\mathrm{Var}}_e(y \mid x) \approx \mathrm{Var}_s\Big(\mathbb{E}_\ell\Big[y^{(s,\ell)}(x)\Big]\Big), \tag{19}$$

element-wise per target. The physical-space aleatoric and epistemic standard deviations (σ_a and σ_e) are then given by the square root of $\widehat{\mathrm{Var}}_a(y \mid x)$ and $\widehat{\mathrm{Var}}_e(y \mid x)$, respectively. During inference, the computational burdens scale with the product of S and L. Therefore, to balance computational costs and statistical accuracy, we set $S = L = 50$ during inference. A sensitivity analysis shows that this results in approximately 400-fold decrease in computation time and less than 4% change in uncertainty when compared to S and L of 1000. More details on sensitivity estimates and inference throughput are detailed in Appendix D and Appendix E.

3.4 Loss Function

Model training minimizes the negative evidence lower bound (ELBO), which balances data fit and model complexity,

$$\mathcal{L}_{\mathrm{ELBO}} = -\mathbb{E}_{\theta\sim q_\phi(\theta)}\left[\log p(\tilde{y} \mid \tilde{x}, \theta)\right] + \mathrm{KL}[q_\phi(\theta) \,\|\, p(\theta)]\,, \tag{20}$$

where the expectation is approximated with Monte Carlo samples from the variational posterior $q_\phi(\theta)$ (Blundell et al., 2015). The first term is the (Gaussian) negative log-likelihood in transformed space, and the second is the KL (Kullback–Leibler divergence) term applied to the Bayesian layers, regularizing the approximate posterior toward the prior.

For a given training pair $(\tilde{x}, \tilde{y})$ and a weight and bias sample $\theta \sim q_\phi(\theta)$, we assume a Gaussian likelihood with diagonal covariance where $\tilde{\mu}_\theta(\tilde{x})$ and $\tilde{\sigma}_{a,\theta}(\tilde{x})$ are the network outputs in latent space,

$$p(\tilde{y} \mid \tilde{x}, \theta) = \mathcal{N}\big(\tilde{y};\, \tilde{\mu}_\theta(\tilde{x}),\, \mathrm{diag}\big(\tilde{\sigma}_{a,\theta}(\tilde{x})^{\odot 2}\big)\big)\,. \tag{21}$$

For training stability, we use a deterministic warm-up of E_{warm} epochs in which the variational posterior over Bayesian-layer parameters is $q_\phi(\theta) = \mathcal{N}(\mu_\theta, \sigma_\theta^2)$ with σ_θ fixed to a near-zero constant $\sigma_{\mathrm{warm}} \approx 0$. In this regime, sampling yields $\theta \approx \mu_\theta$, so Bayesian layers behave nearly deterministically. After warm-up, we allow σ_θ to be learned and initialize it to a small but nonzero value σ_{init}, enabling epistemic uncertainty to develop during subsequent training epochs.

3.5 Model Optimization

We optimized the performance of the hybrid BNN model using a grid search over seven hyperparameters. For each hyperparameter set, we trained the model using 80% of the ADDA data points, 10% for validation during model selection, and held out the remaining 10% for testing; performance was evaluated using the negative log-likelihood (NLL) term of the ELBO loss function. To verify that random splitting did not introduce significant information leakage, we computed test-to-train nearest-neighbor Euclidean distances in standardized input-feature space (excluding the homogeneous sphere inputs); the calculated distances had a minimum of 0.029, a 1st percentile of 0.040, and a median of 0.176, indicating that test samples are generally well separated from their nearest training neighbors. Appendix F provides full details of the hyperparameter search ranges, optimal values, and stopping criteria. Additional training diagnostics and test-set performance plots are provided in Appendix G.

4 Hybrid Bayesian Neural Network Outperforms Homogeneous Sphere and Core-Shell Approximations

We evaluate the hybrid BNN in two complementary ways. First, we assess its capability as a surrogate by comparing its predictions with ADDA reference values on the held-out test set, establishing whether the surrogate can reproduce the underlying morphology-resolving physics. Second, we compare the surrogate against homogeneous-sphere and core-shell approximations, which are the baselines used in large-scale atmospheric models and treat BC-containing particles as homogeneous or concentric volume-equivalent spheres. While the accuracy advantage over these baselines is inherited from ADDA physics, the distinct contributions of the hybrid BNN framework are the computationally efficient access to ADDA-quality optics along with per-particle predictive uncertainty with aleatoric-epistemic decomposition.

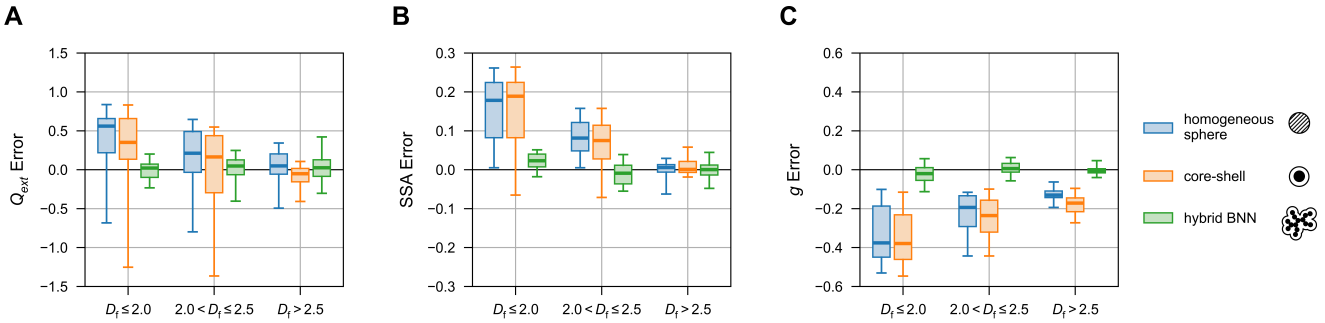


Figure 3. Accuracy of hybrid BNN compared with homogeneous sphere and core-shell approximations for calculating **(A)** extinction efficiency (Q_{ext}), **(B)** single-scattering albedo (SSA), and **(C)** asymmetry parameter (g). The accuracy of each model is quantified for particles with varying fractal dimension (D_{f}). Overall, homogeneous sphere and core-shell approximations have improved accuracy when $D_{\mathrm{f}} > 2.5$, while the hybrid BNN is accurate across the range of D_{f}. Lines represent median values, boxes represent the interquartile range, and error bars represent 5th and 95th percentiles.

We further analyzed the accuracy of each model for predicting the optical properties of BC-containing particles in three groups based on their shape. We quantified particle shape using the fractal dimension (D_{f}), which is distinct from $D_{\mathrm{f,c}}$ because it accounts for the shape of the BC core and coating, as opposed to just the BC core. The structure factor of each coated particle was used to calculate D_{f} (Sorensen, 2001). We separated particles into three groups based on their D_{f}. The first group has $D_{\mathrm{f}} \leq 2.0$ and represents uncoated and thinly coated particles that have not undergone significant core compaction. The second group includes particles with $2.0 < D_{\mathrm{f}} \leq 2.5$, and represents moderately coated uncompacted particles and thinly coated particles that have undergone some degree of core compaction, but their surface morphologies are not spherical. The final group has $D_{\mathrm{f}} > 2.5$, and represents particles with morphologies approaching spherical. These particles include thickly coated particles that have not undergone core compaction and moderately to thickly coated particles which have undergone significant core compaction.

The hybrid BNN was more accurate than homogeneous sphere and core-shell approximations for particles with $D_{\mathrm{f}} \leq 2.5$ (Figure 3a). However, the accuracy of homogeneous sphere and core-shell approximations for predicting Q_{ext} improved as D_{f} increased, a trend which has been demonstrated in previous studies (Schwarz et al., 2008; Lack and Cappa, 2010; Shamjad et al., 2012). The hybrid BNN predicted Q_{ext} of particles with $D_{\mathrm{f}} \leq 2.0$ with median error of 0.02, while homogeneous sphere and core-shell assumptions had median errors of 0.56 and 0.35, respectively. The hybrid BNN also outperformed homogeneous sphere and core-shell approximations when predicting Q_{ext} of particles with $2.0 < D_{\mathrm{f}} \leq 2.5$ (median errors of 0.05, 0.21, and 0.17, respectively). Conversely, the hybrid BNN was comparable to homogeneous sphere and core-shell approximations when predicting Q_{ext} of particles with $D_{\mathrm{f}} > 2.5$ (median errors of 0.03, 0.05, and -0.05, respectively). The hybrid BNN replicates the observational trend of increased accuracy of homogeneous sphere and core-shell approximations for particles with $D_{\mathrm{f}} > 2.5$, but is more accurate than these approximations for BC-containing particles with $D_{\mathrm{f}} \leq 2.5$.

The hybrid BNN also gave the most accurate predictions of SSA for particles with $D_f \leq 2.0$ (Figure 3b). The hybrid BNN predicted SSA with median error of 0.02 for $D_f \leq 2.0$, while homogeneous sphere and core-shell approximations both had median errors > 0.17. The reduced accuracy of Mie theory for BC-containing particles with $D_f \leq 2.0$ is rooted in their complex morphologies. Given that scattering is sensitive to the morphology of particles (Bohren and Huffman, 2008), approximating particles with $D_f \leq 2.0$ as spheres leads to large errors when calculating their scattering properties. The hybrid BNN accurately captured differences in the scattering properties of thinly coated BC-containing particles because it was trained using particles with realistically complex morphologies. The accuracy of homogeneous sphere and core-shell approximations improved for particles with $2.0 < D_f \leq 2.5$ (median errors of 0.08), but the hybrid BNN was still more accurate for these particles (median error of -0.01). Similar to predictions of Q_{ext}, the hybrid BNN was comparable to homogeneous sphere and core-shell approximations for predicting SSA when $D_f > 2.5$ (median errors all < 0.01).

The hybrid BNN was also the most accurate for predicting g, with median error between -0.02 and 0.01. This is a substantial improvement over homogeneous sphere and core-shell approximations, which both systematically underestimate g (median errors between -0.38 and -0.13) (Figure 3c). Unlike predictions of SSA, g was also underestimated by homogeneous sphere and core-shell approximations even for BC-containing particles with $D_f > 2.5$. This is because directional scattering properties are highly sensitive to the internal morphology of particles. Non-spherical aggregates introduce complex anisotropic internal fields and coherence patterns that are not present in symmetric spheres (Bohren and Huffman, 2008). Additionally, adjoining monomers introduce near-field coupling that also affects directional scattering. Incomplete representation of the internal structure of BC-containing particles thus leads to inaccurate estimates of g under homogeneous sphere or core-shell approximations. The lower errors produced by the hybrid BNN relative to ADDA-calculated optical properties show that the hybrid BNN is able to effectively learn the results of the morphology-resolving ADDA physics code without the need to run the more computationally expensive model. Therefore, the advantage of the hybrid BNN model is rooted in its ability to provide mean and uncertainty estimates for optical properties that are reflective of morphology-resolving physics without the need to run an expensive morphology-resolving optics model.

The hybrid BNN outputs also allow for calculation of scattering and absorption efficiencies of particles (Q_{abs} and Q_{sca}). Because BC-containing particles play an important role in total absorption, their Q_{abs} is of particular interest for radiative transfer modeling. However, Q_{abs} is often not supplied directly to radiative transfer models, and is instead given by the product of Q_{ext} and $1-SSA$. Therefore, we have trained the hybrid BNN to predict Q_{ext} and SSA and focused our analysis more heavily on these variables. Nevertheless, the hybrid BNN has similar accuracy to homogeneous sphere and core-shell approximations when Q_{abs} is calculated from Q_{ext} and SSA across particle D_f (see Figure H1). Previous studies have shown that absorption by BC-containing particles is largely insensitive to particle shape (Beeler and Chakrabarty, 2022; Fierce et al., 2020), a trend which is reflected by hybrid BNN predictions. The hybrid BNN was significantly more accurate for predicting Q_{sca} than homogeneous sphere and core-shell approximations, regardless of particle shape. A more detailed comparison of the hybrid BNN model and core-shell and homogeneous sphere approximations for Q_{abs} and Q_{sca} can be found in Appendix H.

Finally, predictions from the hybrid BNN were also compared with those using the Rayleigh-Debye Gans (RDG) approximation (Sorensen et al., 2018). The RDG approximation accounts for the non-spherical shape of BC-containing particles, but

it is not commonly implemented in large-scale atmospheric models. The hybrid BNN was more accurate than the RDG approximation for calculating Q_{ext} and g, and had similar accuracy as the RDG approximation for calculating SSA. More details on the RDG approximation calculations and comparisons with hybrid BNN predictions can be found in Appendix I.

5 Mean and Uncertainty Predictions using Hybrid Bayesian Neural Network

The hybrid BNN model allowed for calculation of the mean optical cross-sections of BC aggregates (with absorbing and non-absorbing coatings) along with their associated uncertainties (Figure 4). We calculated the absorption cross-section (C_{abs}), scattering cross-section (C_{sca}), and asymmetry parameter (g) of particles with various amounts of BC, which is quantified by the core volume equivalent diameter (core d_{ve}). For these calculations, the $D_{\text{f},c}$ that was input to BNN was calculated from V/V_0 following Eqs. (1-3). It should be noted that all particles included in these calculations have $D_{\text{f},c} < 2.5$. The difference between hybrid BNN-calculated optical properties and those calculated using homogeneous sphere or core-shell approximations should be interpreted as the contribution of particle morphology to the optical property of interest. The calculated C_{abs} of BC-containing particles was larger from the hybrid BNN than from homogeneous sphere and core-shell Mie theory for uncoated and thinly coated particles (Figure 4a-c). This trend agrees well with previous work which showed that the absorption efficiency of compact particles is smaller when compared to equivalent fractal particles (Beeler and Chakrabarty, 2022).

Conversely, the predicted C_{sca} of pure BC aggregates ($V/V_0 = 1$) was smaller using the hybrid BNN than for homogeneous sphere and core-shell approximations across core sizes (Figure 4d). The observed lower C_{sca} for uncoated particles is likely due to the optical effects arising from the porosity of fractal BC particles, which has been demonstrated in previous studies (Sorensen, 2001). The C_{sca} of a non-spherical particle is proportional to the effective refractive index (m_{eff}) of a particle, given by

$$\phi\left(\frac{m^2-1}{m^2+2}\right) = \left(\frac{m_{eff}^2-1}{m_{eff}^2+2}\right), \tag{22}$$

where ϕ is the volume packing fraction and m is the refractive index of the material. The value of $|m_{eff}-1|$ is therefore inversely proportional to the porosity of a particle, and higher values of $|m_{eff}-1|$ lead to higher C_{sca}. Aggregates have much higher porosity than equivalent spheres, leading to lower $|m_{eff}-1|$ and lower C_{sca} for aggregates when compared to equivalent spheres. During coating, the ϕ and m_{eff} of particles evolve as coating fills voids within aggregates and compaction occurs. This leads to complex relationships between core d_{ve} and V/V_0 (Figure 4e). Finally, the m_{eff} of BC-containing aggregates approaches that of an equivalent sphere as particles become thickly coated, and relationships between core d_{ve} and C_{sca} approach homogeneous sphere and core-shell approximations (Figure 4f).

The calculated g from the hybrid BNN is larger than those using homogeneous sphere and core-shell approximations regardless of size or coating amount (Figure 4g-i). This trend is also rooted in the sensitivity of directional scattering to the internal morphology of particles, which is not accurately captured by homogeneous sphere and core-shell approximations.

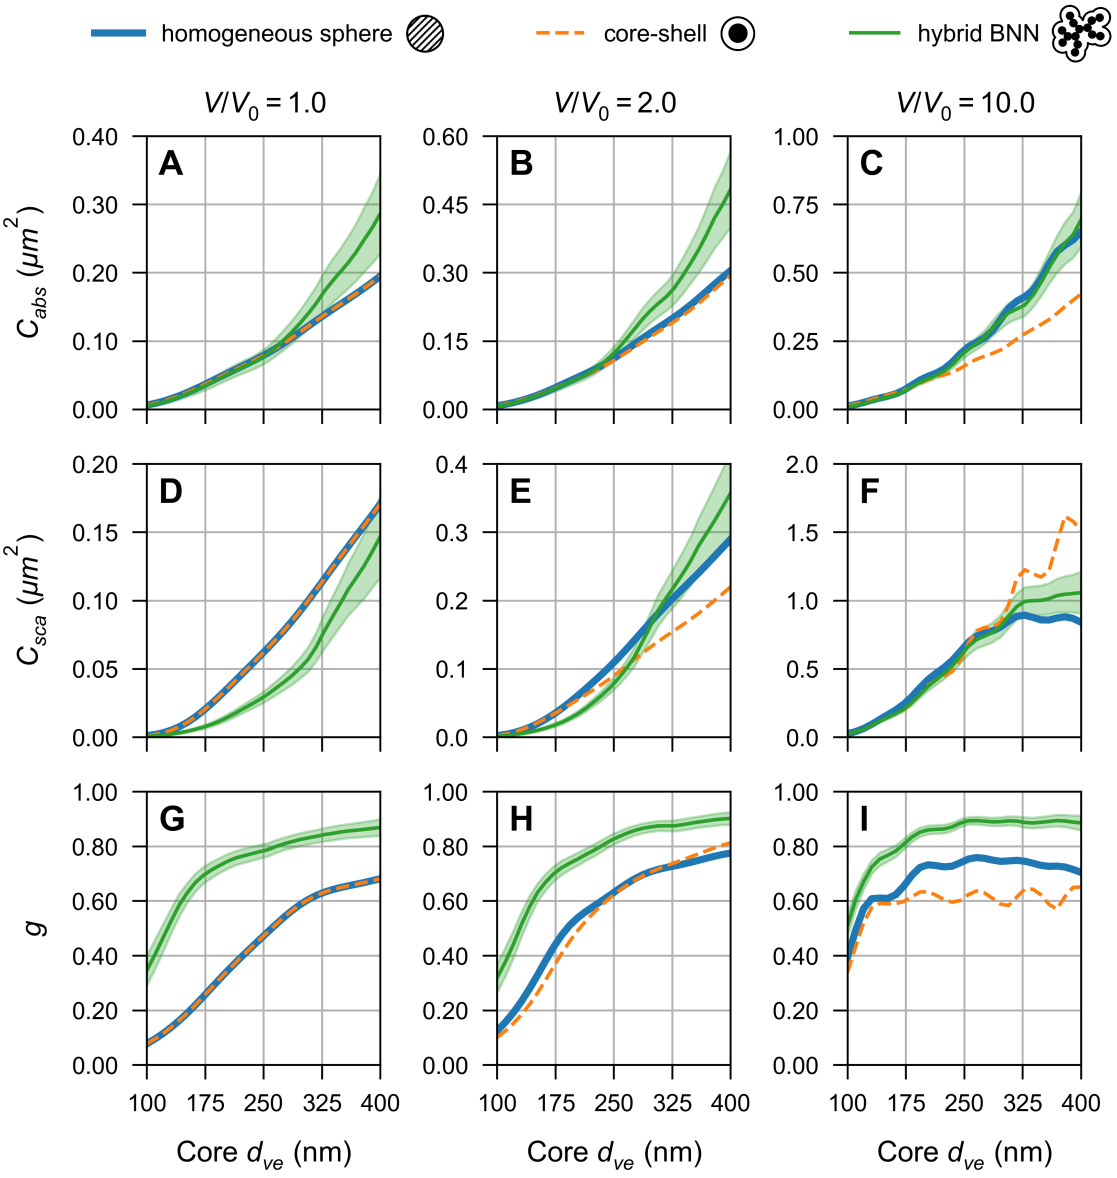


Figure 4. Absorption and scattering predictions at 532 nm using hybrid BNN compared with homogeneous sphere and core-shell approximations. **(A-C)** The absorption cross-section (C_{abs}) of particles predicted by the hybrid BNN is slightly larger than those predicted by the homogeneous sphere and core-shell approximations. **(D-F)** The scattering cross-section (C_{sca}) predicted by the hybrid BNN is smaller than those predicted by the homogeneous sphere and core-shell approximations. **(G-I)** The asymmetry parameter (g) predicted by the hybrid BNN is consistently larger than those predicted by the homogeneous sphere and core-shell approximations. Lines show predicted mean values and shaded areas represent the total standard deviation (computed from the sum of aleatoric and epistemic variances).

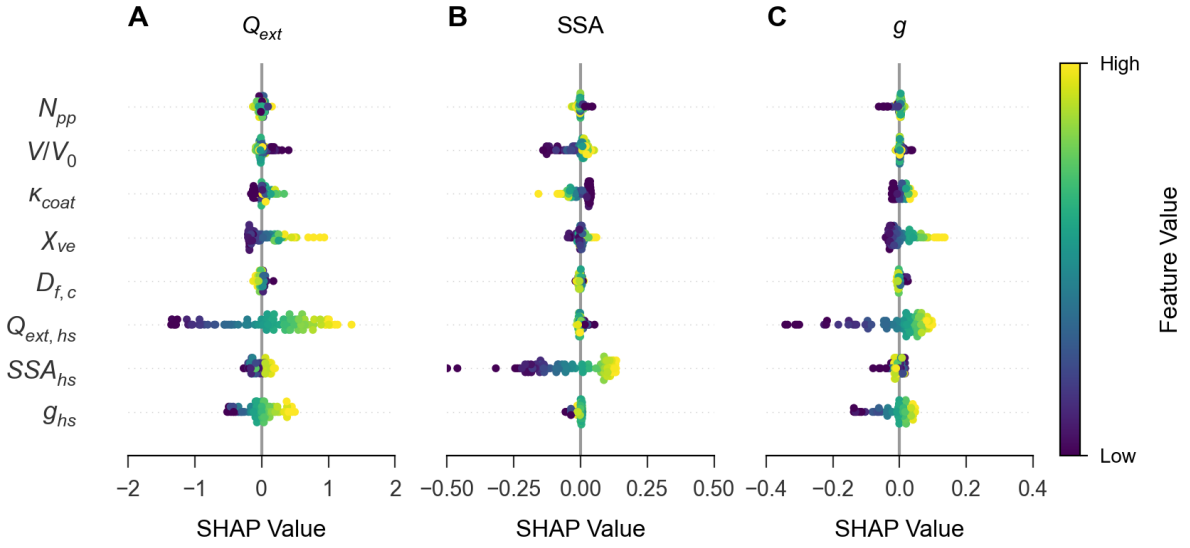


Figure 5. Shapley additive explanation (SHAP) values for extinction efficiency (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) in hybrid BNN. **(A)** Calculated Q_{ext} is most sensitive to χ_{ve} and $Q_{\mathrm{ext,hs}}$ (broad SHAP distributions). **(B)** Calculated SSA is most sensitive to input V/V_0, κ_{coat}, and SSA_{hs}. **(C)** Calculated g is most sensitive to χ_{ve} and $Q_{\mathrm{ext,hs}}$.

6 Input Feature Importance

We calculated Shapley additive explanation (SHAP) values for the eight input features to examine the influence of each input feature on the predicted optical properties using the hybrid BNN model (Figure 5). A broad SHAP value distribution indicates that the given feature has a large influence on the given output, while narrow distributions indicate small influence on the output value (Lundberg and Lee, 2017). High feature values clustered with high SHAP values indicate that increasing the value of the given feature tends to increase the value of the output (and vice versa), while randomly distributed feature values and SHAP values indicate that the feature has a small influence on the output. Our analysis showed that χ_{ve} had a significant effect on Q_{ext}, and that increasing χ_{ve} led to increased Q_{ext} (Figure 5a). This agrees with well-established trends in aerosol optical properties, which state that extinction efficiency increases as particle size parameter increases (Bohren and Huffman, 2008; Kokhanovsky, 2008). The SHAP value distributions also showed that increasing the amount of coating (increasing V/V_0) led to decreased Q_{ext}. This trend is likely rooted in the subtleties of the refractive index of particles. As V/V_0 increases, the real and imaginary parts of the particle refractive index (volume-averaged) decrease. Decreasing the real and imaginary parts of the particle refractive index tends to reduce particle-light interactions and therefore decreases Q_{ext}. However, the exact relationship is not strictly monotonic and depends on other factors such as particle shape and χ_{ve}. Similarly, increasing κ_{coat} while keeping V/V_0 constant increases the absorption efficiency of particles, which in turn increases Q_{ext}.

We also found that V/V_0 had a large influence on SSA (Figure 5b), but that the effect of V/V_0 on SSA diminished at high V/V_0. This trend agrees with previous laboratory studies which show that the addition of scattering material leads to an

increase in SSA (Kalbermatter et al., 2021; Xu et al., 2018). The diminishing effect of V/V_0 on SSA at high V/V_0 is also in good agreement with our current understanding of particle optics. As V/V_0 increases, the extinction efficiency of particles becomes dominated by scattering while absorption effectively saturates (Beeler and Chakrabarty, 2022), both of which lead to plateauing SSA. However, increasing κ_{coat} increases the contribution of absorption to total extinction, leading to decreased SSA, an effect that was also well represented in the SHAP value distributions.

Finally, SHAP distributions show that χ_{ve} had a large effect on g (Figure 5c). This reflects fundamental principles in directional light scattering, which predict that forward scattering is enhanced as the size of particles relative to the wavelength of incident light increases (Bohren and Huffman, 2008; Kokhanovsky, 2008). The coating absorptivity also had a significant effect on g, with increasing κ_{coat} leading to increased g. This trend also agrees well with fundamental particle optics, which predict that when χ_{ve} is < 1, increasing κ_{coat} effectively decreases large-angle scattering while leaving small-angle scattering relatively unchanged. This leads to increased g when κ_{coat} increases and $\chi_{ve} < 1$. However, complex coherence and interference patterns when $\chi_{ve} > 1$ lead to more complex relationships between g and κ_{coat}.

The input homogeneous sphere optical descriptors ($Q_{\mathrm{ext,hs}}$, SSA_{hs}, and g_{hs}) also had a significant impact on all outputs of the hybrid BNN. For instance, a large input value for $Q_{\mathrm{ext,hs}}$ led to a larger predicted Q_{ext} value. To evaluate the relative importance of the combined microphysical descriptors with the homogeneous sphere optical descriptors on the target variables, we calculated the total group importance, given by the mean of $\sum_{j \in G} |\psi_j|$, where j is the feature in group G and ψ_j are the corresponding SHAP values. The total group importance gives the cumulative influence of the group on the prediction. This analysis showed that the total importance of the microphysical descriptors was 0.38, 0.08, and 0.07 for Q_{ext}, SSA, and g, respectively. The importance of the homogeneous sphere optical descriptors on the predictions was 0.80, 0.13, and 0.10 for Q_{ext}, SSA, and g, respectively. These findings collectively show that the homogeneous sphere optical descriptors are very important for hybrid BNN predictions. Including the homogeneous sphere optical properties as an input provides the hybrid BNN with a large fraction of the underlying physics involved in light scattering and absorption, which allowed the hybrid BNN to be trained on a relatively small number of data points. The importance of $Q_{\mathrm{ext,hs}}$, SSA_{hs}, and g_{hs} for hybrid BNN prediction also indicates that the hybrid BNN may be accurate for refractive indices and monomer sizes that were not represented in the training set if the inputs are reflective of the new refractive index or monomer size.

7 Extrapolation of Model Predictions Across Refractive Indices and Monomer Sizes

The hybrid BNN training set included optical properties of BC-containing particles with monomer radii of 20 nm, coatings with real refractive indices ($Re(\mathrm{RI_{coat}})$) of 1.6, and BC with refractive index ($\mathrm{RI_{BC}}$) of $1.95 + 0.79i$. We performed a perturbed-parameter ensemble to evaluate the accuracy of the hybrid BNN for particles with monomer radii, $Re(\mathrm{RI_{coat}})$, and $\mathrm{RI_{BC}}$ outside of these values. We randomly selected 10 points from the held-out test set and systematically changed the input monomer radii in ADDA simulations. The same process was repeated for simulations in which $Re(\mathrm{RI_{coat}})$ and $\mathrm{RI_{BC}}$ were altered, resulting in a total of 90 additional ADDA simulations. We then used the BNN to predict the optical properties of the particles and compared the BNN predictions to the ADDA simulations. We benchmarked model error across $Re(\mathrm{RI_{coat}})$, $\mathrm{RI_{BC}}$, and monomer radii

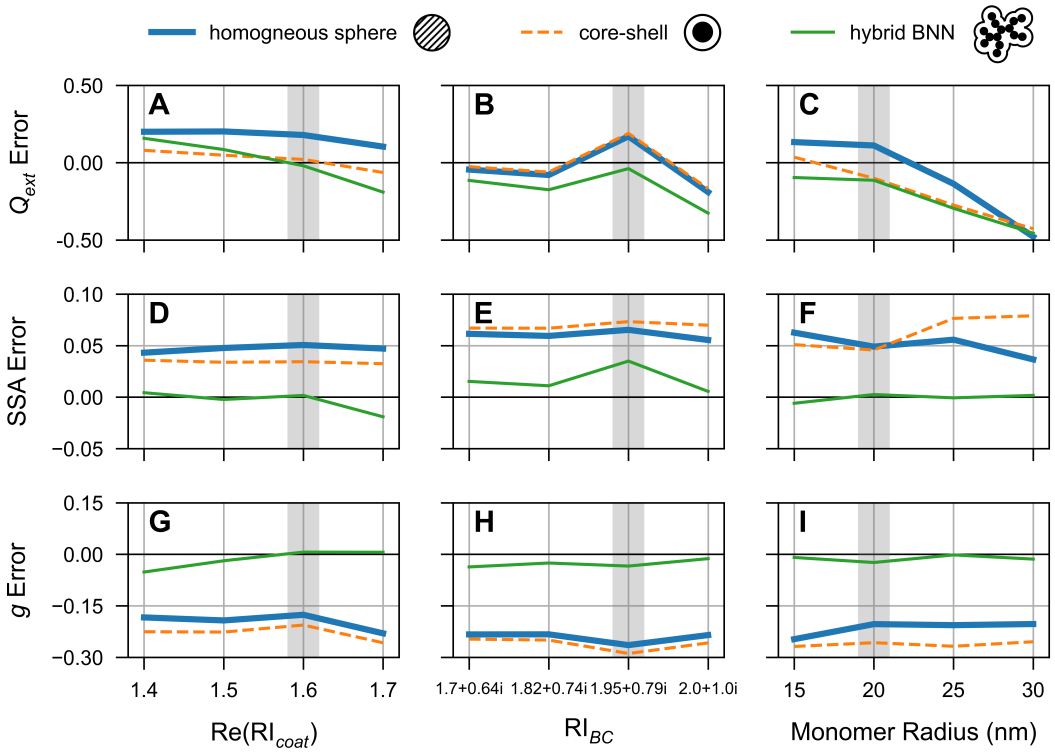


Figure 6. Mean error in extinction efficiency (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) calculated by the hybrid BNN and using homogeneous sphere and core-shell approximations for a perturbed parameter ensemble. Here, error is given by the difference between predicted and true values. **(A-C)** The hybrid BNN shows varying accuracy as the real part of the coating refractive index ($Re(\mathrm{RI_{coat}})$), BC refractive index ($\mathrm{RI_{BC}}$), or monomer radius deviate from training set values. However, it still has comparable accuracy to homogeneous sphere and core-shell approximations. **(D-I)** The hybrid BNN is more accurate than homogeneous sphere and core-shell approximations when calculating SSA and g across all $Re(\mathrm{RI_{coat}})$, $\mathrm{RI_{BC}}$, and monomer radius. Shaded areas in each panel indicate values that were used in the training set and lines show average values.

using the average difference between calculated and true values, where true values are given by ADDA simulations (Figure 6). The shaded regions in Figure 6 denote the parameter values represented in training.

Deviations in $Re(\mathrm{RI_{coat}})$, $\mathrm{RI_{BC}}$, or monomer radius from training values affected the accuracy of the hybrid BNN for predicting Q_{ext}, but the hybrid BNN still had comparable errors to homogeneous sphere and core-shell approximations (Figure 6a-c). These results indicate that the hybrid BNN loses predictive accuracy for calculating Q_{ext} under high relative humidities, where the refractive index of water would contribute significantly to $Re(\mathrm{RI_{coat}})$. However, it is likely that the hybrid BNN would still be more accurate than homogeneous sphere or core-shell approximations under dry to moderate relative humidities, assuming dry coatings have $Re(\mathrm{RI_{coat}})$ between 1.5 and 1.6 (Bond and Bergstrom, 2006). The hybrid BNN was more accurate

than homogeneous sphere and core-shell approximations for calculating SSA and g across $Re(\mathrm{RI_{coat}})$, $\mathrm{RI_{BC}}$, and monomer radii (Figure 6d-i).

It should be noted that the monomer radius, $Re(\mathrm{RI_{coat}})$, and $\mathrm{RI_{BC}}$ are not directly included as inputs to the hybrid BNN. However, changing the monomer radius alters χ_{ve}, and changing the refractive indices of BC or the coating material changes the Mie theory-calculated optical properties ($Q_{\mathrm{ext,hs}}$, SSA_{hs}, and g_{hs}), both of which are used as inputs to the hybrid BNN. Including $Q_{\mathrm{ext,hs}}$, SSA_{hs}, and g_{hs} as hybrid BNN inputs allows for extrapolation of BNN predictions to monomer sizes and refractive indices that were not used to train the model with little loss in accuracy.

8 Future Experiment Design Using Epistemic Uncertainty Predictions

One of the unique features of the hybrid BNN is the ability to separate aleatoric and epistemic uncertainty in predicted optical properties. To demonstrate that epistemic uncertainty can identify sparse regions in input-parameter space, the hybrid BNN was retrained after masking approximately $12\,\%$ of the training data in the region $70 \leq N_{\mathrm{pp}} \leq 180$ and $8 \leq V/V_0 \leq 25$. Test particles within this masked region were retained for evaluation but were not used during re-training. Epistemic uncertainty increased more inside the masked region than outside it for all three target variables (see Appendix J). Because epistemic uncertainty increases in regions with sparse training data, patterns in epistemic uncertainty can be used to identify regions of the parameter space that would most benefit from additional training data. Therefore, uncertainty estimates from the hybrid BNN can be leveraged to guide the design of future simulations that would most effectively improve model performance.

We used the hybrid BNN to predict aleatoric and epistemic uncertainty in Q_{ext}, SSA, and g at three wavelengths commonly used in optical instruments (405 nm, 532 nm, and 880 nm). Here, we quantified error in optical properties as a percentage of the predicted mean value. Overall, we found that aleatoric uncertainty is highest for smaller, thinly coated aggregates, which is due to the high variability in the morphology of these particles as they compact due to the accumulation of coatings (see Appendix K). We also found that epistemic Q_{ext} uncertainty at 405 nm is highest for large, thickly coated particles with $N_{\mathrm{pp}} > 800$ and $V/V_0 > 40$ (Figure 7a). Epistemic Q_{ext} uncertainty at 532 nm is highest for thinly coated particles ($V/V_0 < 5$) and $N_{\mathrm{pp}} < 20$ or $N_{\mathrm{pp}} > 700$ (Figure 7b). At 880 nm, epistemic Q_{ext} uncertainty is highest for small, thinly coated particles with $N_{\mathrm{pp}} < 50$ and $V/V_0 < 10$ (Figure 7c). For SSA and g, epistemic uncertainty across wavelengths is largest for small, thinly coated particles with $N_{\mathrm{pp}} < 100$ and $V/V_0 < 10$ (Figure 7d-i).

Given that ADDA simulations are extremely computationally expensive, this analysis can be used to identify regions of the input parameter space that are either sparsely represented or where there is significant variability in the optical properties of particles. Including additional simulations that fall within these areas of the input parameter space would be the most beneficial for improved model performance. By targeting these sparsely represented input parameters, future simulations can be designed more efficiently, maximizing model improvement while reducing computational burdens. These results also show that future laboratory experiments with BC-containing particles with $N_{\mathrm{pp}} < 100$ and $V/V_0 < 10$ would also aid in constraining model uncertainty. Previous studies have shown that particles in this region are in the midst of moderate restructuring that is dependent on N_{pp} (Beeler et al., 2025). Therefore, these particles are expected to have more morphological variability than particles with

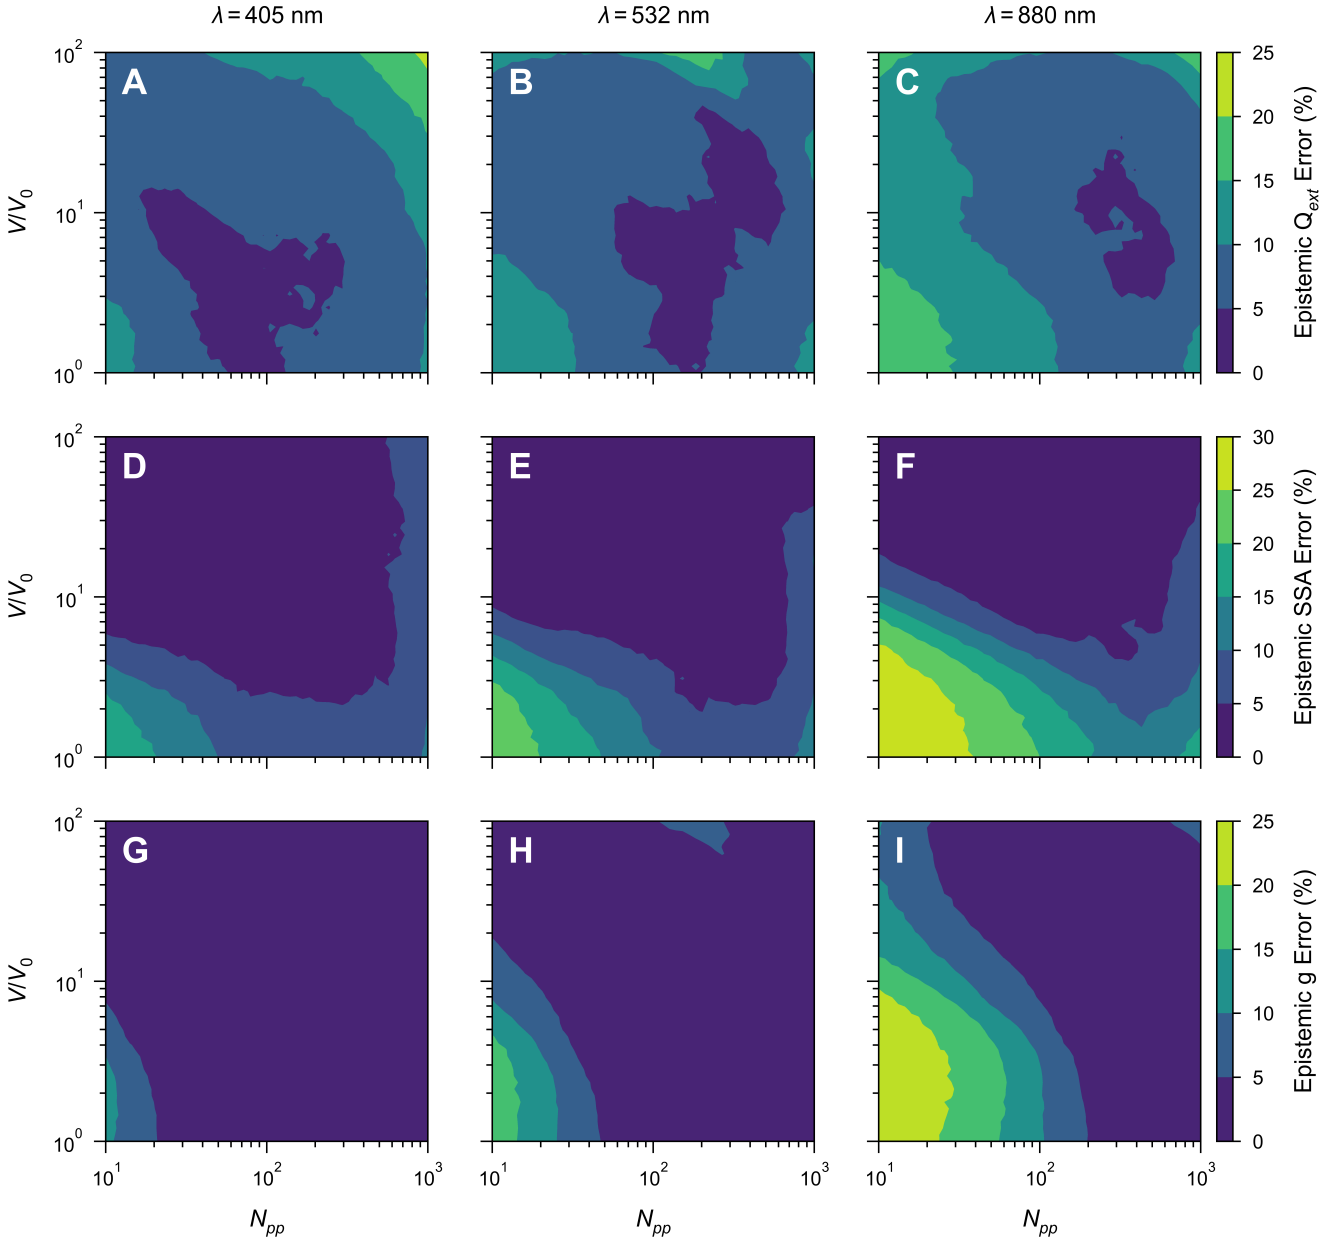


Figure 7. Epistemic uncertainty in predicted optical properties can be used for targeted future experiments aimed at reducing model uncertainty. **(A-C)** Epistemic uncertainty in predicted Q_{ext} is highest for large, thickly coated particles ($N_{\mathrm{pp}} > 800$ and $V/V_0 > 40$) at 405 nm. Epistemic uncertainty for SSA and g is largest for small, thinly coated particles at 532 and 880 nm. **(D-I)** For SSA and g, epistemic uncertainty across wavelengths is largest for small, thinly coated particles with $N_{\mathrm{pp}} < 100$ and $V/V_0 < 10$.

$N_{pp} > 100$ and $V/V_0 > 10$. Additional laboratory experiments targeting these particles would therefore be particularly valuable for constraining their optical properties. Measurements that jointly constrain aggregate size, coating amount, core and whole-particle morphology, and wavelength-resolved optical properties would provide especially useful constraints for evaluating and informing future versions of the model.

9 Discussion and Conclusions

Efficiently predicting the optical properties of BC-containing particles remains a longstanding challenge in atmospheric models and remote sensing algorithms (Ramanathan and Carmichael, 2008; Bond and Bergstrom, 2006). Traditional atmospheric models employ simplifications to BC morphology and mixing with other aerosol components to reduce computational burdens (Liu et al., 2016). However, these simplifications often do not capture the complexity of light scattering and absorption by irregularly shaped BC-containing particles. Previous studies have attempted to improve predictions of the optical properties of BC-containing particles using machine learning-based approaches (Qin et al., 2025; Romshoo et al., 2024; Luo et al., 2018). However, these studies use single-particle optics models which do not accurately capture near-field interactions between adjacent monomers in BC aggregates and do not quantify uncertainty in model predictions due to variability in particle morphology.

We developed a hybrid BNN to predict the optical properties of BC-containing particles and their associated uncertainties. The model was trained using ADDA simulations of single-particle optical properties, which accurately represent the morphology and near-field interactions between monomers in BC-containing particles (Yurkin and Hoekstra, 2011). The hybrid BNN was able to effectively learn the morphology-resolving physics of ADDA simulations while reducing computation time by 2-6 orders of magnitude. The hybrid BNN takes in five inputs that are easily tracked or derived in large-scale models and fully describe the size, shape, and composition of BC-containing particles. The hybrid BNN also uses the optical properties of a volume-equivalent homogeneous sphere as an input feature. The optical properties of a volume-equivalent homogeneous sphere can be calculated with low computational cost and inform the hybrid BNN of the complex physics involved in particle-light interactions. Our results showed that the hybrid BNN more accurately predicts mean BC optical properties than traditional homogeneous-sphere or core–shell approximations, with the largest improvements for SSA and g. The hybrid BNN also predicts uncertainty in the predicted optical properties of BC-containing particles.

We found that the size of the particle relative to the wavelength, imaginary part of the coating refractive index, and the optical properties of a homogeneous volume-equivalent sphere were the most influential input features for model predictions. We also showed that the hybrid BNN is able to extrapolate to monomer sizes, BC refractive indices, and coating refractive indices that were not represented in the training set with little loss in accuracy. This is because inclusion of homogeneous sphere optical descriptors encodes much of the physics of these parameters on target variables, and therefore makes hybrid BNN predictions applicable to various assumed BC and coating refractive indices and monomer sizes. Finally, uncertainty decomposition in the hybrid BNN allows for identification of regions in the input parameter space that are underrepresented in the training dataset

or have significant variability in the optical properties. This allows for targeted future additions to the training dataset that will be most effective in improving model predictions.

It should be noted that the hybrid BNN introduced here is a single-particle model that is not suitable for direct incorporation into large-scale atmospheric models. Multiple lines of evidence show that large-scale models do not accurately represent the mixing state of BC-containing particles, which is extremely important for predicting their optical properties (Fierce et al., 2020, 2016). The hybrid BNN introduced here does not address this gap, and therefore is not likely to improve BC optics calculations in large-scale models. However, our hybrid BNN can be used as a computationally-efficient surrogate model for calculating BC optics in particle-resolved models, which often contain thousands of particles. Bulk optical properties from particle-resolved simulations can then be used to evaluate optics schemes in large-scale models, which will be the subject of future studies. Additionally, the hybrid BNN framework could be extended to predict scattering quantities relevant for remote sensing applications, such as phase function or depolarization. This will also be the subject of future studies.

Overall, our results present the first machine learning-based model for BC optical properties with uncertainty quantification. It is also the first machine learning-based model for BC optical properties that is trained using ADDA simulations, to our knowledge. Our hybrid BNN accurately predicts the optical properties of BC-containing particles across the morphology-composition space. Quantification of uncertainty in single-particle optical properties allows for propagation to uncertainty in bulk optical properties, which are important for comparison against measurements and for probabilistic predictions of BC behavior in large-scale atmospheric models. Improved predictions of BC optical properties will be a major step forward for improved representation of anthropogenic and biomass burning emissions in large-scale atmospheric models.

Code and data availability. The DDA-BNN codebase, including the training pipeline, training dataset, and inference tools, is publicly available at https://github.com/pnnl/DDA-BNN under the BSD 2-Clause License. The exact version of the code used to produce the results in this paper is archived on Zenodo at https://doi.org/10.5281/zenodo.21478183 (Beeler et al., 2026b). The training dataset is available within the GitHub repository and is also archived on Zenodo at https://doi.org/10.5281/zenodo.19324185 (Beeler et al., 2026a).

Appendix A: ADDA Run Configuration

Version 1.5.0-alpha2 of the ADDA code was used to calculate the optical properties of simulated BC-containing particles. The ADDA user manual recommends keeping the number of dipoles per wavelength (λ/d) above 10 to keep errors in calculated optical properties to within a few percent. Since computation time increases steeply with the number of dipoles (Figure A1a), particles were discretized using a dipole length of 0.0033 μm to balance simulation accuracy with computational cost ($\lambda/d > 90$ for all simulations). ADDA simulations were performed on the Deception high-performance computing cluster at Pacific Northwest National Laboratory using a single AMD EPYC 7502 CPU core for each simulation. Finally, we calculated the optical properties of particles for 150 randomly generated orientations and used the average optical properties over the 150 orientations for hybrid BNN training and testing. We randomly selected 40 particles from the hybrid BNN test set and calculated

their optical properties for 300 random orientations. We found that using the average optical properties over 150 orientations deviated by 0.1% or less compared to the average value over 300 orientations (Figure A1b). We therefore expect errors in the calculated optical properties from ADDA due to discretization and orientation-averaging to be negligible.

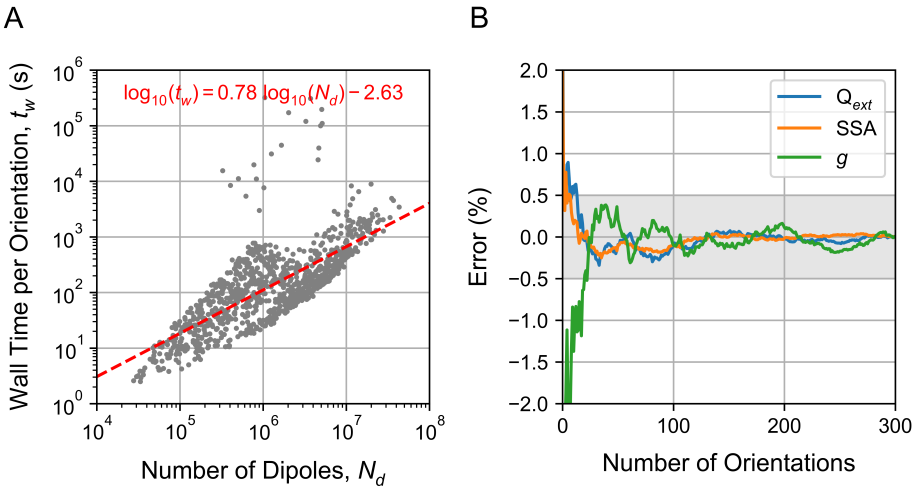


Figure A1. Timing and errors from ADDA simulations. (A) Computation time per orientation scales approximately as a power law with the number of dipoles. (B) Optical properties converge after approximately 100 orientations.

Appendix B: Z-Score Normalization

Z-score normalization ensures that each variable has mean zero and standard deviation one in the normalized space. The forward Z-score transformation is

$$Z(x) = \frac{x - \mu_x}{\sigma_x}, \tag{B1}$$

for input variables, and

$$Z(y_{tf}) = \frac{y_{tf} - \mu_{y_{tf}}}{\sigma_{y_{tf}}}, \tag{B2}$$

for transformed target variables (Fei et al., 2021). Here, μ and σ are the mean and standard deviation of the corresponding variable computed on the training set only; the same training-set statistics are then used to normalize the validation and test sets. The input and transformed target variables are mapped and stored as $x \xrightarrow{Z} \tilde{x}$ and $y \xrightarrow{T} y_{tf} \xrightarrow{Z} \tilde{y}$.

Appendix C: Uncertainty Scaling Calibration

We applied post-hoc *uncertainty scaling* to calibrate the predictive uncertainty of the hybrid BNN for each target (Q_{ext}, SSA, and g). Specifically, we fit a per-target multiplicative scaling factor τ on the validation set by minimizing the Gaussian negative

log-likelihood in latent space (Section 3.3.2). Table C1 reports the fitted τ values and the empirical coverage of nominal 90% prediction intervals before (Pre) and after (Post) calibration on the validation and test sets. Values of $\tau < 1$ indicate that the raw BNN uncertainties were overestimated. Figure C1 shows reliability diagrams comparing expected versus observed coverage; curves closer to the 1:1 line indicate better-calibrated predictive uncertainty.

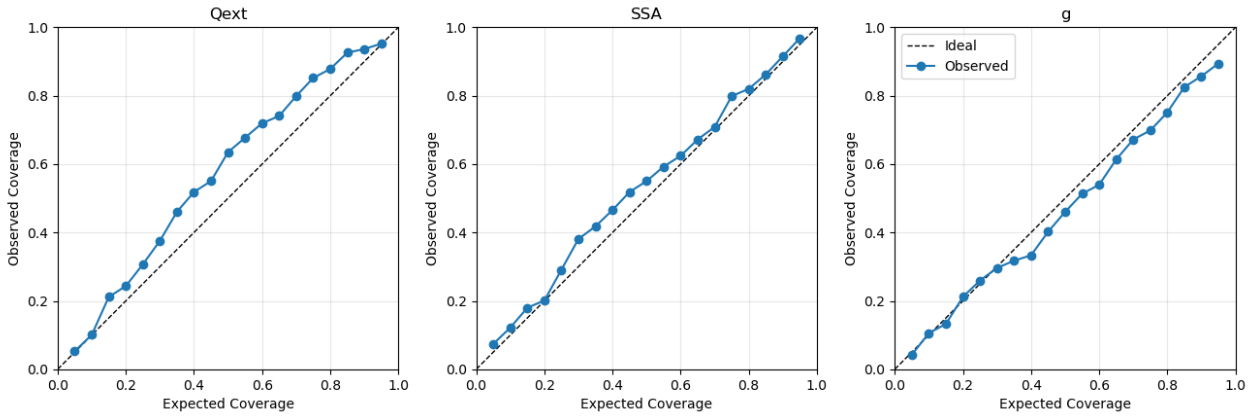


Figure C1. Reliability diagrams for hybrid BNN predictive uncertainty for Q_{ext}, SSA, and g. The dashed line indicates ideal calibration (observed coverage equals expected coverage), and the points/solid line show observed coverage. Deviations from the 1:1 line indicate miscalibration.

Table C1. Uncertainty scaling calibration results. All fitted values have $\tau < 1$, indicating that the raw BNN overestimated uncertainty.

		Validation		Test	
Target	τ	Pre	Post	Pre	Post
Q_{ext}	0.521	98.4%	92.0%	97.9%	93.7%
SSA	0.440	100%	88.3%	98.9%	90.5%
g	0.530	100%	87.8%	98.4%	85.2%

Appendix D: Uncertainty Convergence

Predictive uncertainties reported in the main text are computed by nested Monte Carlo sampling, with S posterior parameter draws and L aleatoric samples each. We characterize the convergence of the two components with respect to S and L using σ_e^* and σ_a^* to denote the epistemic and aleatoric standard deviations at a converged reference $(S, L) = (1000, 1000)$. All deviations in this section are expressed with respect to that reference.

Figure D1 shows the two components as functions of S (left column, with L pinned at 1000) and L (right column, with S pinned at 1000), for each of the three targets. As the paper operates at $S = L = 50$ on both axes simultaneously, we additionally report the joint deviation at that point along with several representative combinations in Table D1. Values are pooled-median deviations across 10 test particles and 3 seeds. At the paper default, all six deviations remain within $\pm 4\,\%$.

The choice of $S = L = 50$ is a cost–accuracy trade-off. Nested-MC sampling has $O(S \cdot L)$ cost in both compute and memory, so moving from $(S, L) = (50, 50)$ to the converged reference at $(1000, 1000)$ increases per-particle wall-clock and memory footprint by a factor of ~ 400 while tightening the reported uncertainties by less than $4\,\%$ on any single target. The paper default therefore trades a small, quantified loss of precision for two orders of magnitude of throughput. Absolute inference timings and batch scaling are reported in Appendix E.

Table D1. Deviation of predictive uncertainty at representative (S, L) combinations with respect to the converged reference at $(S, L) = (1000, 1000)$.

		σ_e deviation (%)			σ_a deviation (%)		
S	L	Q_{ext}	SSA	g	Q_{ext}	SSA	g
50	50	−3.9	+3.0	−0.7	−1.3	−2.8	+0.6
10	1000	−17.0	+6.5	−5.1	+2.2	+0.1	−0.2
100	1000	−2.0	−0.9	+4.0	−1.0	+0.2	+0.3
1000	10	+5.0	+3.1	+6.1	−5.5	−5.3	−5.0
1000	100	+0.3	+0.9	+0.8	−0.4	−0.2	−0.6
1000	1000	0	0	0	0	0	0

Appendix E: Inference Throughput

We report the wall-clock cost of the hybrid BNN at the paper default parameters of $S = L = 50$. All timings were measured on a single Intel Core i7-11850H workstation (8 physical cores, 16 GB RAM, CPU-only PyTorch 2.4 with Pyro 1.9) using the median wall clock of 20 repeated inference passes after a single warm-up call. Sensitivity of the reported uncertainties to the (S, L) sampling choice is documented separately in Appendix D.

Single-particle inference (batch $= 1$) is dominated by fixed per-call overhead. As batch size grows the overhead is distributed across particles and the per-particle cost drops toward the asymptotic compute cost of the network forward passes and nested-MC sampling. The per-particle time plateaus around ~ 1 ms for batches of 10^3–10^4 particles, delivering approximately 10^3 particles/s on a single CPU core.

We note that the current inference pipeline is CPU-only, as the Bayesian sampling stack (Pyro's variational inference) does not straightforwardly support CUDA in our configuration. A GPU-enabled implementation would be expected to sustain substantially higher throughput at large batch sizes; extending the pipeline to GPU is left as future engineering work.

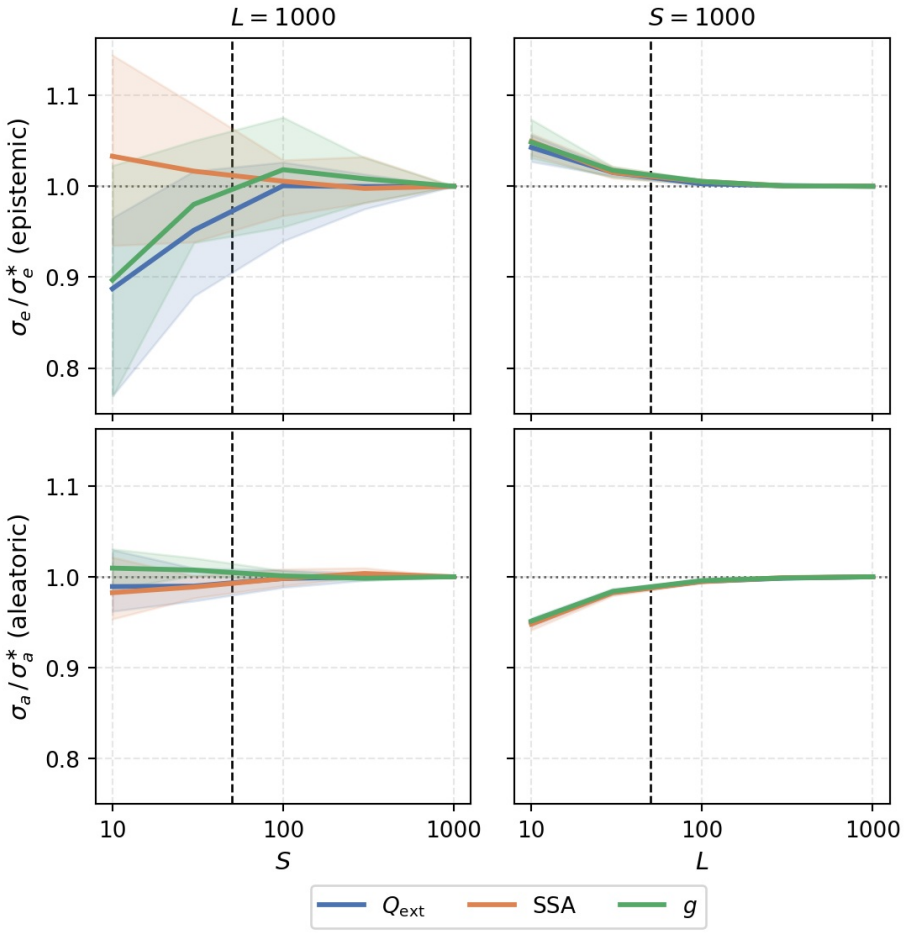


Figure D1. Monte Carlo convergence of predictive uncertainty. Curves show the pooled median across 10 test particles and 3 seeds of the ratio σ/σ^*, where σ^* is the value at the converged reference $(S, L) = (1000, 1000)$. Shaded bands show the interquartile range across the pooled sample; the dashed vertical line marks the paper default $S = L = 50$.

Table E1. Hybrid BNN inference cost as a function of batch size at $S = L = 50$. Per-particle time and throughput are computed from wall clock and batch size.

Batch size	Wall (s)	Per-particle (ms)	Throughput (particles/s)
1	0.15	151	7
10	0.15	14.6	69
100	0.22	2.2	455
1000	1.04	1.0	961
10000	11.0	1.1	908

Appendix F: Hybrid BNN Hyperparameter Selection

Table F1 summarizes the hyperparameters explored during model selection and the optimal configuration. *Layers* denotes the total number of hidden layers in the network, and *Bayesian layers* denotes the number of variational (Bayesian) layers used to represent weight uncertainty. *Bayesian layer position* indicates where the Bayesian layer(s) were placed within the network (middle or last). The *learning-rate scale factor* multiplies the base learning rate used by the optimizer for any Bayesian layers (post E_{warm}), and the *prior standard deviation* sets the scale of the Gaussian prior on Bayesian-layer weights.

The warm-up parameter E_{warm} denotes the number of initial epochs used for the deterministic warm-up described in the main text. Training used early stopping based on the validation loss: if the validation loss did not decrease by at least $\Delta_{\min} = 0.001$ for 100 consecutive epochs, training was terminated.

For the optimal model, the hidden-layer widths (all using ReLU activations) were 128, 128, 128 (Bayesian), 64, and 32 units. Additional training diagnostics and test-set performance plots corresponding to the selected hyperparameter configuration are provided in Appendix G.

Table F1. Hyperparameters explored during hybrid BNN model selection (Section 3.5), including sampled values and the optimal configuration.

Hyperparameter	Sampled Values	Optimal Value
Layers	5, 7, 9	5
Bayesian layers	1, 2, 3	1
Bayesian layer position	middle, last	middle
Learning rate	1.0×10^{-5}, 3.0×10^{-5}, 3.0×10^{-4}	3.0×10^{-5}
Learning rate scale factor	0.1, 0.5, 1.0	0.1
Prior standard deviation	0.1, 0.5, 1.0, 2.0	1.0
E_{warm}	0, 50, 100	50

Appendix G: Hybrid BNN Training and Test-Set Performance

Figure G1 shows the evolution of the negative log-likelihood (NLL) and KL-divergence components of the hybrid BNN objective during training for the training and validation sets. The similar trajectories of training and validation NLL indicate stable optimization and limited overfitting over the training period.

Figure G2 compares test-set predictions from the hybrid BNN (points with 90% predictive confidence intervals) and the homogeneous-sphere approximation (orange markers) against ADDA reference values (dashed 1:1 line). Points closer to the 1:1 line indicate higher accuracy relative to ADDA, and the BNN error bars summarize predictive uncertainty for individual test samples.

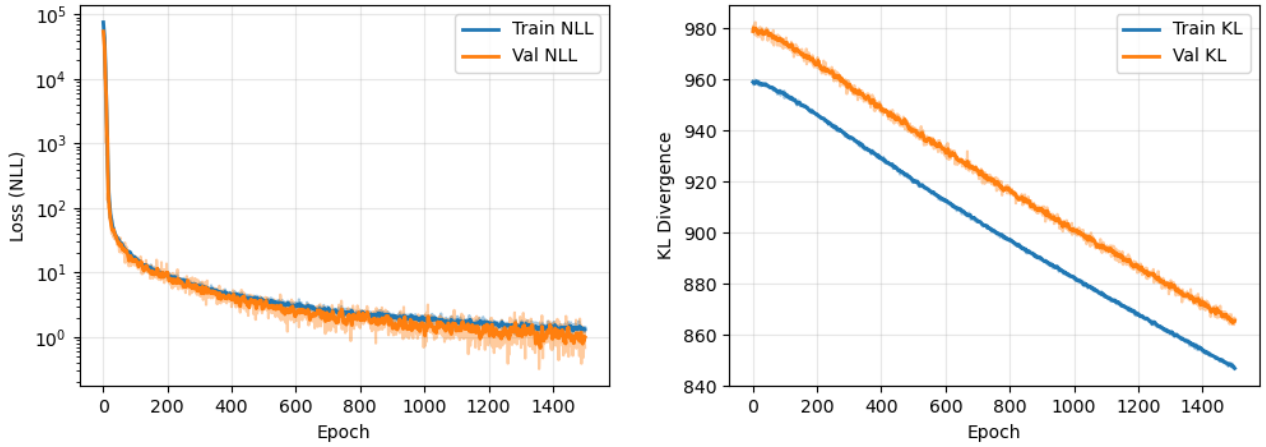


Figure G1. Training history for the hybrid BNN loss components. Left: negative log-likelihood (NLL) for training and validation sets. Right: KL divergence term for training and validation sets.

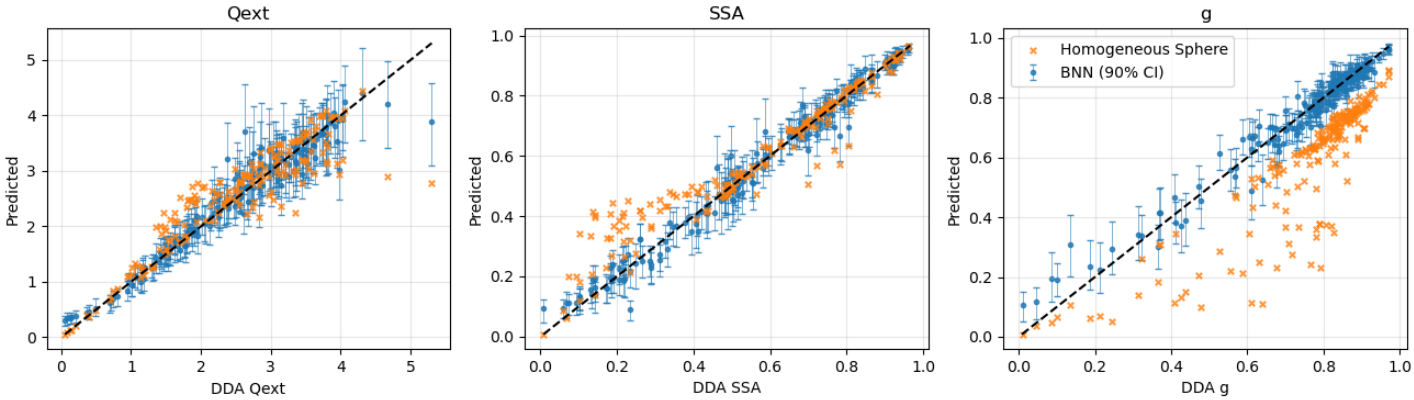


Figure G2. Test-set comparison of predicted versus ADDA reference values for Q_{ext}, SSA, and g. Orange markers show the homogeneous-sphere approximation. Blue points show hybrid BNN predictions with 90% predictive confidence intervals. The dashed line indicates the 1:1 relationship with ADDA.

Appendix H: Absorption and Scattering Efficiencies from Hybrid Bayesian Neural Network

Predictions of Q_{ext} and SSA were combined to calculate the absorption (Q_{abs}) and scattering (Q_{sca}) efficiencies of particles (see Figure H1). The hybrid BNN calculated Q_{abs} with median error of 0.01, 0.03, and 0.01 for $D_{\text{f}} \leq 2.0$, $2.0 < D_{\text{f}} \leq 2.5$, and $D_{\text{f}} > 2.5$, respectively. The BNN was slightly more accurate than homogeneous sphere and core-shell approximations for $D_{\text{f}} \leq 2.0$ and $D_{\text{f}} > 2.5$, but was less accurate for $2.0 < D_{\text{f}} \leq 2.5$. This shows that homogeneous sphere and core-shell approximations can predict Q_{abs} with reasonable accuracy, but the hybrid BNN model also predicts Q_{abs} with similar accuracy. Conversely, the hybrid BNN greatly outperforms homogeneous sphere and core-shell approximations when calculating Q_{sca}. Median error for Q_{sca} is $-0.02 - 0.05$ using the hybrid BNN model, $0.06 - 0.49$ using the homogeneous sphere approximation, and $-0.02 - 0.51$ using the core-shell approximation. Core-shell and homogeneous sphere approximations again have improved accuracy as D_{f} increases. Overall, homogeneous sphere and core-shell approximations may be adequate for calculating absorption, but are not accurate for calculating the scattering properties of BC-containing particles with $D_{\text{f}} \leq 2.5$.

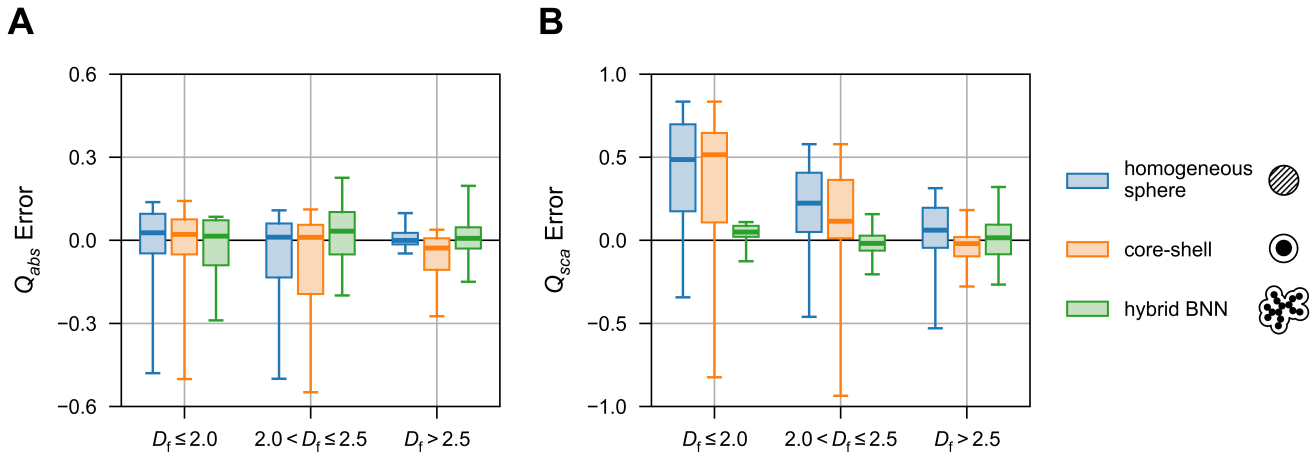


Figure H1. Accuracy of hybrid BNN, homogeneous sphere approximation, and core-shell approximation for calculating absorption efficiency (Q_{abs}) and scattering efficiency (Q_{sca}) as a function of particle shape (quantified by D_{f}). **(A)** The three methods have comparable accuracy for predicting Q_{abs} across particle shapes. **(B)** The hybrid BNN is more accurate than homogeneous sphere and core-shell approximations when predicting Q_{sca}, regardless of particle shape. Bars show median values and errors are 25th and 75th percentiles.

Appendix I: Comparison of Hybrid BNN Predictions with the RDG Approximation

We compared the accuracy of hybrid BNN predictions with those using the Rayleigh-Debye Gans (RDG) approximation. The RDG approximation accounts for the non-spherical shape of BC-containing aggregates by assuming that the monomers comprising a BC-containing particle are within the Rayleigh optics limit (Sorensen et al., 2018). The optical properties of particles in the test set were calculated using Eqs. 16−23 of Sorensen et al. (2018). To account for the effects of coating on RDG

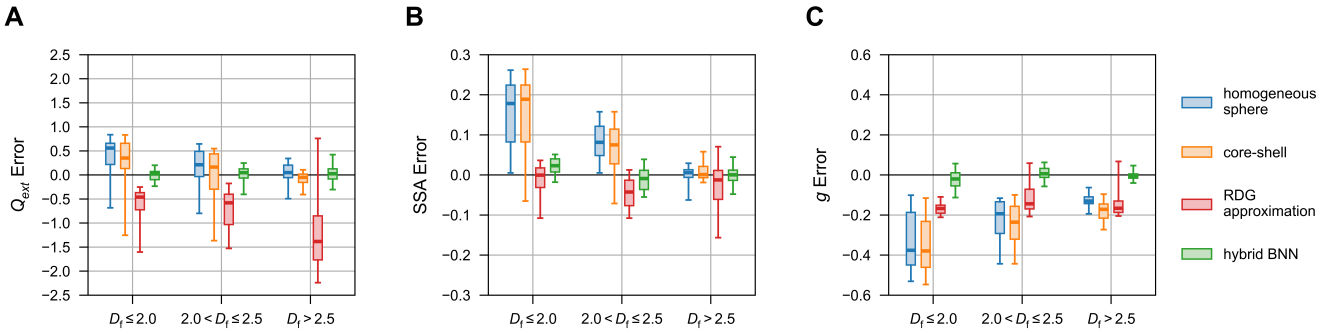


Figure I1. Comparison of hybrid BNN accuracy with the core-shell RDG approximation. The hybrid BNN was more accurate than the RDG approximation for Q_{ext} and g, and had similar accuracy as the RDG approximation for SSA.

approximation calculations, the optical response of each coated monomer is represented using the electrostatic polarizability of a concentric core-shell sphere (β_{cs}). Defining the relative permittivity of the core, shell, and surrounding medium as ε_{c}, ε_{s}, and ε_{m}, respectively, and the core volume fraction (f) as

$$f = \left(\frac{a_{\text{c}}}{a_{\text{o}}}\right)^3, \tag{I1}$$

where a_{c} and a_{o} are the core and outer radii. The polarizability is then calculated using

$$\beta_{\text{cs}} = \frac{(\varepsilon_{\text{s}} - \varepsilon_{\text{m}})(\varepsilon_{\text{c}} + 2\varepsilon_{\text{s}}) + f(\varepsilon_{\text{c}} - \varepsilon_{\text{s}})(\varepsilon_{\text{m}} + 2\varepsilon_{\text{s}})}{(\varepsilon_{\text{s}} + 2\varepsilon_{\text{m}})(\varepsilon_{\text{c}} + 2\varepsilon_{\text{s}}) + 2f(\varepsilon_{\text{s}} - \varepsilon_{\text{m}})(\varepsilon_{\text{c}} - \varepsilon_{\text{s}})}, \tag{I2}$$

with

$$\varepsilon_i = \left(\frac{m_i}{m_{\text{m}}}\right)^2. \tag{I3}$$

Here, m_i denotes the complex refractive index of the core or shell and m_{m} is the refractive index of the surrounding medium ($1.0+0.00i$). This expression reduces to the homogeneous-sphere Lorentz–Lorenz factor when the core and shell have identical optical properties. We therefore use β_{cs} as opposed to the Lorentz-Lorenz factor when calculating $F(m)$ and $E(m)$. It should also be noted that the fractal dimension and fractal prefactor were assumed fixed at 1.78 and 1.35, respectively, in RDG approximation calculations.

The hybrid BNN is more accurate than the RDG approximation for Q_{ext}, with median errors between $(0.02, 0.05)$ and $(-1.4, 0.46)$, respectively. The hybrid BNN had similar accuracy as the RDG approximation for calculating SSA, with median errors between $(-0.01, 0.02)$ and $(-0.04, -0.01)$, respectively. The accuracy of RDG approximation for calculating SSA but not Q_{ext} indicates that RDG approximation tends to underestimate total absorption and/or scattering cross-sections of coated aggregates, but approximately captures their relative contributions to total extinction. Finally, the hybrid BNN is also more accurate than the RDG approximation when calculating g for coated aggregates, with median errors between $(-0.02, 0.01)$ and $(-0.20, -0.15)$, respectively.

Appendix J: Sensitivity of Predictive Uncertainty to Data Sparsity

A defining feature of the hybrid BNN is that its epistemic uncertainty σ_e represents the model's own uncertainty in underrepresented regions during training. As a direct test of that interpretation, we retrained the hybrid BNN after masking approximately $12\,\%$ of the training data in the region $70 \leq N_{\mathrm{pp}} \leq 180$ and $8 \leq V/V_0 \leq 25$. Test particles within this masked region were retained for evaluation but were not used during retraining. All other hyperparameters and the random seed were identical to the baseline model.

Figure J1 shows the per-particle changes in epistemic uncertainty (σ_e) and aleatoric uncertainty (σ_a) between the baseline model (trained on the full dataset) and the retrained model (trained on the masked dataset). For test particles inside the masked region, epistemic uncertainty increased by $14.4\,\%$, $28.6\,\%$, and $11.9\,\%$ for Q_{ext}, SSA, and g, respectively. For test particles outside the masked region, the corresponding increases were $3.9\,\%$, $5.0\,\%$, and $-1.4\,\%$. 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. For test particles inside the masked region, aleatoric uncertainty increased by $6.9\,\%$, $27.9\,\%$, and $3.6\,\%$ for Q_{ext}, SSA, and g, respectively, compared with changes of $0.8\,\%$, $8.7\,\%$, and $-0.4\,\%$ outside the masked 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 epistemic uncertainty responds to reduced training coverage.

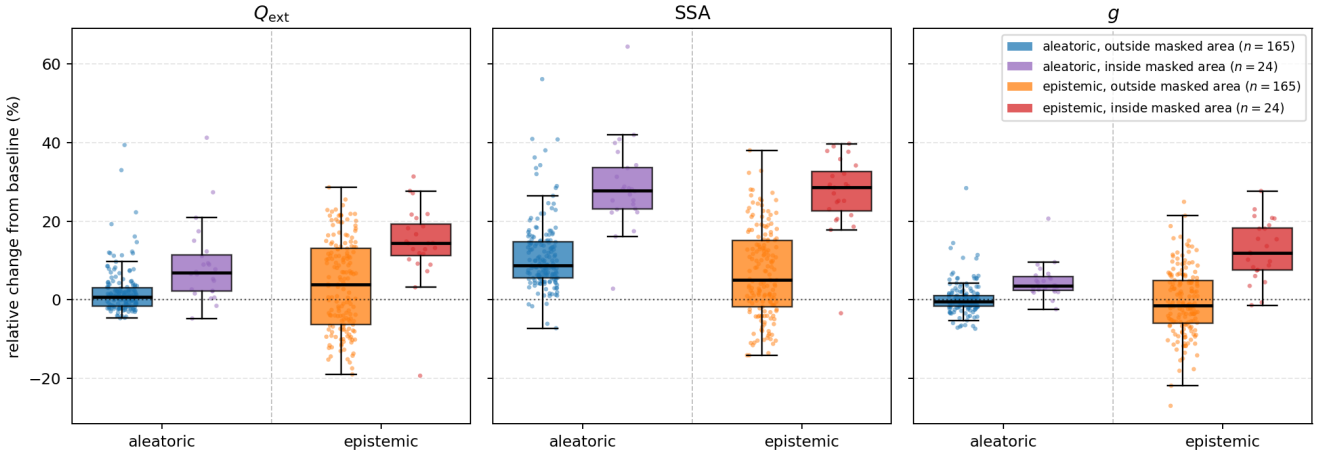


Figure J1. Per-particle changes in epistemic uncertainty (σ_e) and aleatoric uncertainty (σ_a) when the hybrid BNN is retrained after masking approximately $12\,\%$ of the training data in the region $70 \leq N_{\mathrm{pp}} \leq 180$ and $8 \leq V/V_0 \leq 25$. Points show the relative change in each uncertainty component with respect to the baseline model for each held-out test particle.

Appendix K: Aleatoric Uncertainty Estimates

We used the hybrid BNN to predict aleatoric uncertainty in Q_{ext}, SSA, and g at 405 nm, 532 nm, and 880 nm. Errors in optical properties are expressed as a percentage of the predicted mean value. We found that aleatoric uncertainty is highest in regions where particles have the highest morphological variability (uncoated or thinly coated particles). This shows that aleatoric uncertainty is able to identify regions of the parameter space where uncertainty due to particle variability is highest.

Author contributions. **PB**: Conceptualization, Methodology, Software, Data curation, Writing - original draft. **SD**: Methodology, Software, Data curation, Writing - original draft. **LF**: Conceptualization, Methodology, Data curation, Writing - original draft.

Corresponding author: all correspondence to LF (laura.fierce@pnnl.gov).

Competing interests. The authors declare that they have no competing interests.

Acknowledgements. This research was supported by the U.S. Department of Energy (DOE) Laboratory Directed Research and Development program (Linus Pauling Distinguished Postdoctoral Fellowship Program) at Pacific Northwest National Laboratory. L. Fierce was supported by the U.S. Department of Energy (DOE) Atmospheric System Research (ASR) program via the Integrated Cloud, Land-Surface, and Aerosol System Study (ICLASS) Science Focus Area. Pacific Northwest National Laboratory is operated for DOE by Battelle Memorial Institute under contract DE-AC05-76RL01830.

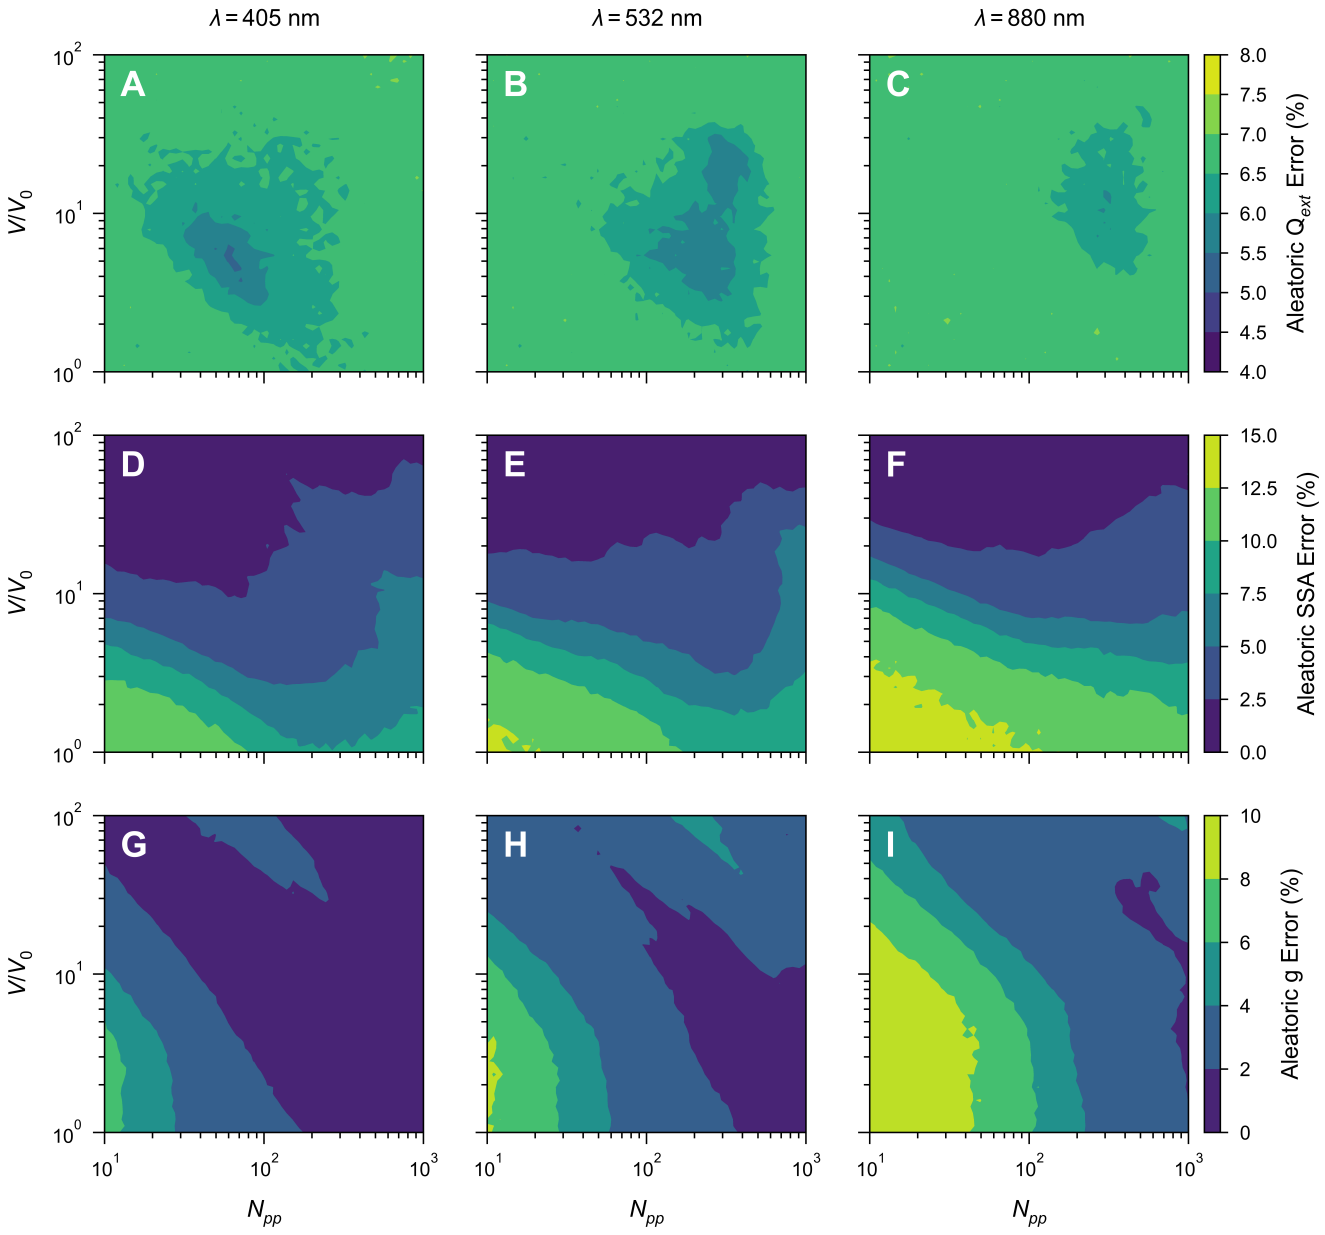


Figure K1. Aleatoric uncertainty as a function of BC core size (quantified by N_{pp}) and coating amount (quantified by V/V_0) at three wavelengths commonly used in aerosol optical instruments. We found that aleatoric uncertainty is generally higher for thinly coated aggregates, which reflects high morphological variability of fractal BC aggregates before coating-induced compaction.

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.

Beeler, P., Donald, S., and Fierce, L.: Training dataset for DDA-BNN v1.0, https://doi.org/10.5281/zenodo.19324185, 2026a.

Beeler, P., Donald, S., and Fierce, L.: DDA-BNN v1.0: A Morphology-Aware Surrogate Model for the Optical Properties of Black Carbon–Containing Particles, https://doi.org/10.5281/zenodo.21478183, 2026b.

Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D.: Weight uncertainty in neural network, in: International conference on machine learning, pp. 1613–1622, PMLR, 2015.

Bohren, C. F. and Huffman, D. R.: Absorption and scattering of light by small particles, John Wiley & Sons, 2008.

Bond, T. C. and Bergstrom, R. W.: Light absorption by carbonaceous particles: An investigative review, Aerosol Science and Technology, 40, 27–67, 2006.

Bond, T. C., Doherty, S. J., Fahey, D. W., Forster, P. M., Berntsen, T., DeAngelo, B. J., Flanner, M. G., Ghan, S., Kärcher, B., and Koch, D.: Bounding the role of black carbon in the climate system: A scientific assessment, Journal of Geophysical Research: Atmospheres, 118, 5380–5552, 2013.

Bonnet, D., Hirtzlin, T., Majumdar, A., Dalgaty, T., Esmanhotto, E., Meli, V., Castellani, N., Martin, S., Nodin, J.-F., and Bourgeois, G.: Bringing uncertainty quantification to the extreme-edge with memristor-based Bayesian neural networks, Nature Communications, 14, 7530, 2023.

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.

China, S., Mazzoleni, C., Gorkowski, K., Aiken, A. C., and Dubey, M. K.: Morphology and mixing state of individual freshly emitted wildfire carbonaceous particles, Nature Communications, 4, 2122, 2013.

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.

Fei, N., Gao, Y., Lu, Z., and Xiang, T.: Z-score normalization, hubness, and few-shot learning, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 142–151, 2021.

Fierce, L., Bond, T. C., Bauer, S. E., Mena, F., and Riemer, N.: Black carbon absorption at the global scale is affected by particle-scale diversity in composition, Nature Communications, 7, 12 361, 2016.

Fierce, L., Onasch, T. B., Cappa, C. D., Mazzoleni, C., China, S., Bhandari, J., Davidovits, P., Fischer, D. A., Helgestad, T., and Lambe, A. T.: Radiative absorption enhancements by black carbon controlled by particle-to-particle heterogeneity in composition, Proceedings of the National Academy of Sciences, 117, 5196–5203, 2020.

Gal, Y. and Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning, Proceedings of the 33rd International Conference on Machine Learning, pp. 1050–1059, 2016.

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.

Kalbermatter, D. M., Močnik, G., Drinovec, L., Visser, B., Röhrbein, J., Oscity, M., Weingartner, E., Hyvärinen, A.-P., and Vasilatou, K.: Response of black carbon and aerosol absorption measuring instruments to laboratory-generated soot coated with controlled amounts of secondary organic matter, Atmospheric Measurement Techniques Discussions, 2021, 1–21, 2021.

Kendall, A. and Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision?, Advances in neural information processing systems, 30, 2017.

Koeylue, U., Xing, Y., and Rosner, D. E.: Fractal morphology analysis of combustion-generated aggregates using angular light scattering and electron microscope images, Langmuir, 11, 4848–4854, 1995.

Kokhanovsky, A. A.: Aerosol optics: light absorption and scattering by particles in the atmosphere, Springer, 2008.

Koshelev, M.: Standard statistical transformations (logarithm and logit) are uniquely determined by the corresponding symmetries, Journal of Uncertain Systems, 9, 103–112, 2015.

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.

Laskin, A., Laskin, J., and Nizkorodov, S. A.: Chemistry of atmospheric brown carbon, Chemical reviews, 115, 4335–4382, 2015.

Li, H. and May, A. A.: Estimating mass-absorption cross-section of ambient black carbon aerosols: Theoretical, empirical, and machine learning models, Aerosol Science and Technology, 56, 980–997, 2022.

Li, J., Zhang, F., Liu, J., Li, W., Wu, K., Hu, S., and Lin, H.: Parameterization of optical properties for liquid cloud droplets containing black carbon based on neural network, Optics Express, 31, 40 124–40 141, 2023.

Liu, C., Li, J., Yin, Y., Zhu, B., and Feng, Q.: Optical properties of black carbon aggregates with non-absorptive coating, Journal of Quantitative Spectroscopy and Radiative Transfer, 187, 443–452, 2017a.

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, 2017b.

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.

Liu, S., Aiken, A. C., Gorkowski, K., Dubey, M. K., Cappa, C. D., Williams, L. R., Herndon, S. C., Massoli, P., Fortner, E. C., and Chhabra, P. S.: Enhanced light absorption by mixed source black and brown carbon particles in UK winter, Nature Communications, 6, 8435, 2015.

Liu, X., Ma, P.-L., Wang, H., Tilmes, S., Singh, B., Easter, R., Ghan, S., and Rasch, P.: Description and evaluation of a new four-mode version of the Modal Aerosol Module (MAM4) within version 5.3 of the Community Atmosphere Model, Geoscientific Model Development, 9, 505–522, 2016.

Lundberg, S. M. and Lee, S.-I.: A unified approach to interpreting model predictions, Advances in neural information processing systems, 30, 2017.

Luo, J., Zhang, Y., Wang, F., Wang, J., and Zhang, Q.: Applying machine learning to estimate the optical properties of black carbon fractal aggregates, Journal of Quantitative Spectroscopy and Radiative Transfer, 215, 1–8, 2018.

Olivier, A., Shields, M. D., and Graham-Brady, L.: Bayesian neural networks for uncertainty quantification in data-driven materials modeling, Computer Methods in Applied Mechanics and Engineering, 386, 114 079, 2021.

Qin, Z., Wu, J., Wang, H., Zhang, Y., and Zhang, Q.: Using convolutional neural networks to predict the optical properties of coated black carbon, Journal of Quantitative Spectroscopy and Radiative Transfer, 333, 109 326, 2025.

Ramanathan, V. and Carmichael, G.: Global and regional climate changes due to black carbon, Nature Geoscience, 1, 221–227, 2008.

Romshoo, B., Patil, J., Michels, T., Müller, T., Kloft, M., and Pöhlker, M.: Improving the predictions of black carbon (BC) optical properties at various aging stages using a machine-learning-based approach, Atmospheric Chemistry and Physics, 24, 8821–8846, 2024.

Saleh, R., Cheng, Z., and Atwi, K.: The brown–black continuum of light-absorbing combustion aerosols, Environmental Science & Technology Letters, 5, 508–513, 2018.

Samset, B. H., Sand, M., Smith, C. J., Bauer, S. E., Forster, P. M., Fuglestvedt, J. S., Osprey, S., and Schleussner, C.-F.: Climate impacts from a removal of anthropogenic aerosol emissions, Geophysical Research Letters, 45, 1020–1029, 2018.

Schwarz, J. P., Spackman, J., Fahey, D., Gao, R., Lohmann, U., Stier, P., Watts, L., Thomson, D., Lack, D., Pfister, L., et al.: Coatings and their enhancement of black carbon light absorption in the tropical atmosphere, Journal of Geophysical Research: Atmospheres, 113, 2008.

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.

Sorensen, C. M., Yon, J., Liu, F., Maughan, J., Heinson, W. R., and Berg, M. J.: Light scattering and absorption by fractal aggregates including soot, Journal of Quantitative Spectroscopy and Radiative Transfer, 217, 459–473, 2018.

Sumlin, B. J., Heinson, W. R., and Chakrabarty, R. K.: Retrieving the aerosol complex refractive index using PyMieScatt: A Mie computational package with visualization capabilities, Journal of Quantitative Spectroscopy and Radiative Transfer, 205, 127–134, 2018.

Xu, X., Zhao, W., Qian, X., Wang, S., Fang, B., Zhang, Q., Zhang, W., Venables, D. S., Chen, W., Huang, Y., et al.: The influence of photochemical aging on light absorption of atmospheric black carbon and aerosol single-scattering albedo, Atmospheric Chemistry and Physics, 18, 16 829–16 844, 2018.

Yurkin, M. A. and Hoekstra, A. G.: The discrete-dipole-approximation code ADDA: Capabilities and known limitations, Journal of Quantitative Spectroscopy and Radiative Transfer, 112, 2234–2247, 2011.

Zhao, G., Zhao, C., Kuang, Y., Bian, Y., Tao, J., Shen, C., and Yu, Y.: Calculating the aerosol asymmetry factor based on measurements from the humidified nephelometer system, Atmospheric Chemistry and Physics, 18, 9049–9060, 2018.