



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

Payton Beeler¹, Sam Donald², and Laura Fierce¹

¹Atmospheric, Climate, and Earth Sciences Division, Pacific Northwest National Laboratory, Richland, WA, 99354, USA

²Physical and Computational Sciences, Pacific Northwest National Laboratory, Richland, WA, 99354, USA

Correspondence: Laura Fierce (laura.fierce@pnl.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 numerically exact 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
25 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). BC-containing particles are emitted as frac-
tal 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 can restructure aggregates toward more compact con-
30 figurations 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). 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).

To achieve computational efficiency, large-scale atmospheric and radiative transfer models typically represent BC-containing
35 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 mor-
phologies and compositions (Yurkin and Hoekstra, 2011), but are too computationally expensive for routine use in regional
40 or global modeling and are impractical for embedding directly into iterative workflows such as inverse modeling, data as-
similation, 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
45 et al., 2017). 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 par-
50 ticles, 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 pa-
rameter (g) and associated uncertainties for individual particles, allowing for the decomposition of total model uncertainty into
55 aleatoric and epistemic components. We benchmark performance against homogeneous-sphere and core-shell Mie approxima-
tions 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.



<p><u>Ensemble I</u></p> <p>990 particles $r(V/V_0, D_{f,c}) = 0$</p>	<p>$V/V_0 = 3.0$ $D_{f,c} = 1.8$</p> 	<p>$V/V_0 = 1.0$ $D_{f,c} = 3.0$</p> 	<p>$V/V_0 = 1.5$ $D_{f,c} = 2.5$</p> 
<p><u>Ensemble II</u></p> <p>990 particles $r(V/V_0, D_{f,c}) > 0$</p>	<p>$V/V_0 = 1.0$ $D_{f,c} = 1.8$</p> 	<p>$V/V_0 = 2.6$ $D_{f,c} = 1.9$</p> 	<p>$V/V_0 = 14.0$ $D_{f,c} = 2.2$</p> 

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 uncorrelated total-to-black carbon volume ratio (V/V_0) and core fractal dimension ($D_{f,c}$). Therefore, the correlation coefficient (r) between $D_{f,c}$ and V/V_0 is zero. The particles in Ensemble II (bottom row) represent compacting particles and have V/V_0 that is calculated based on $D_{f,c}$ following Beeler et al. (2025). These particles therefore have corrected V/V_0 and $D_{f,c}$ ($r > 0$).

2 Training Set Generation

60 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

65 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.



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 ($D_{f,c}$), 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 (RI_{BC}), the coating refractive index at 550 nm ($RI_{coat,550\text{ nm}}$), and the absorption Ångstrom exponent of the coating (AAE_{coat}).

Ensemble	Data Points	N_{pp}	r_{pp} (nm)	V/V_0	$D_{f,c}$	λ (nm)	RI_{BC}	$RI_{coat,550\text{ nm}}$	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 – 2500	$1.95 + 0.79i$	$1.6 + [0.00 – 0.25]i$	0.35 – 7.6

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

70 Non-compacting BC-containing particles are taken from a previous publication (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). 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
75 primary particles with radii of 20 nm. The degree of compaction is quantified using the core fractal dimension ($D_{f,c}$), with fresh aggregates having low $D_{f,c}$ (1.8) and compact aggregates having high $D_{f,c}$ (> 2.5) (Koeylue et al., 1995). It is important to note that $D_{f,c}$ and V/V_0 are not correlated in Ensemble I. We simulated 900 particles in Ensemble I, with N_{pp} ranging from 13 – 1189, $D_{f,c}$ ranging from 1.8 – 3.0, and V/V_0 ranging from 1 – 75 (see Table 1). The optical properties are calculated at
80 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).

2.2 Ensemble II: Compacting Black Carbon-Containing Particles

Previous studies have developed scaling laws for $D_{f,c}$ as a function of V/V_0 using detailed laboratory measurements and single-particle simulations (Beeler et al., 2025). Therefore, the $D_{f,c}$ of particles in this 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
85 to 1000, and the V/V_0 of each particle is randomly selected between 1 and 100 (see Table 1). The $D_{f,c}$ of each particle is then given by

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



where x_c is given by

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

90 and x_f is given by

$$x_f = 1.07x_0 N_{pp}^{-0.026}. \quad (3)$$

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} , V/V_0 , and $D_{f,c}$ are generated using diffusion limited cluster-cluster aggregation, Van-der Waals compaction, and capillary coating models that are described in detail in Beeler et al. (2025). Ensemble II
95 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 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

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

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

$$AAE_{\text{coat}} = 7.6 \exp[-12.4\kappa_{550 \text{ nm}}]. \quad (5)$$

These relationships represent 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 2500 nm,
105 the refractive index of BC is assumed constant at $1.95 + 0.79i$, the radius of primary particles is 20 nm, $\kappa_{550 \text{ nm}}$ 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 Numerically Exact Optical Properties of Simulated Black Carbon-Containing Particles

We use the Amsterdam discrete dipole approximation 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
110 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. The optical properties of particles in Ensemble I are calculated at 532, 880, 1200, or 2500 nm. We calculate the optical properties of Ensemble II particles at wavelengths randomly distributed amongst 50 nm-spaced bins ranging from 300 to 2500 nm.



115 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 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 coefficient (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) as the target variables of the hybrid BNN model. 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 BC-containing particle properties as input variables to the hybrid BNN model: N_{pp} , V/V_0 , $D_{f,c}$, 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 ($D_{f,c}$), 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 $(0, 1)$ (e.g., 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 $g < 0$

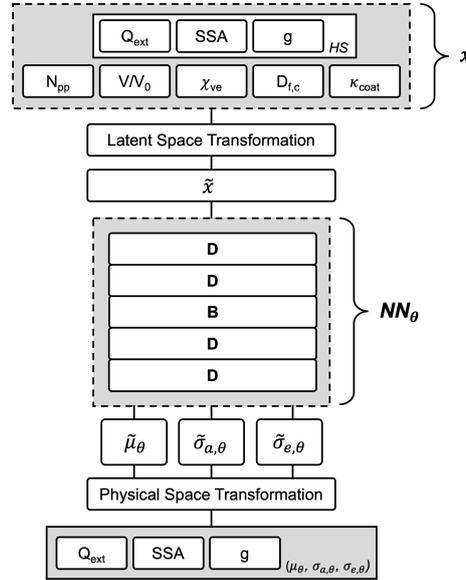


Figure 2. Example of the hybrid Bayesian neural network architecture. Inputs are standardized using Z-score normalization, and target variables are mapped to an unbounded latent space using the transform T followed by Z-score normalization. The neural network (NN_θ) consists of a combination of deterministic and Bayesian layers producing a latent-space mean ($\tilde{\mu}_\theta$), aleatoric uncertainty ($\tilde{\sigma}_{\alpha,\theta}$), and epistemic uncertainty ($\tilde{\sigma}_{\epsilon,\theta}$). 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.

145 (Zhao et al., 2018). Therefore, we use the logit transform for g . Specifically, T is given by:

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

150 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 A.

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:

$$155 \quad NN_\theta = [D, B, D], \quad (7)$$



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

$$z \mapsto \phi(Wz + b), \quad (8)$$

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.

The outputs of the model, for a given set of sampled weights 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 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. 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 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,

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

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

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

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

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



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

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

3.3.2 Uncertainty Scaling

Although Eqs. (10–11) provide an aleatoric–epistemic decomposition in latent space, the overall predictive uncertainty can be miscalibrated leading to prediction intervals that are systematically too wide or narrow. We therefore apply per-target
190 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 ($\tilde{\sigma}_{\text{tot},\theta}(\tilde{x})$), and the validation targets transformed into latent space ($\tilde{y}_t = ZT(y_t)$). For each target dimension j , we solve

$$195 \quad \tau_j = \arg \min_{\tau > 0} \frac{1}{N} \sum_{i=1}^N \text{NLL}(\tilde{y}_{t,j}(x_i); \tilde{\mu}_{\theta,j}(\tilde{x}_i), \tau \tilde{\sigma}_{\text{tot},\theta,j}(\tilde{x}_i)). \quad (13)$$

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 B.

3.3.3 Propagation to Physical Space

200 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,

$$205 \quad \tilde{\mu}_\tau^{(s)}(\tilde{x}) = \tilde{\mu}_\theta(\tilde{x}) + \tau \odot \left(\tilde{\mu}^{(s)}(\tilde{x}) - \tilde{\mu}_\theta(\tilde{x}) \right). \quad (14)$$

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

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

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

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

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

$$y^{(s,\ell)}(x) = T^{-1} \left(Z^{-1}(\tilde{y}^{(s,\ell)}(\tilde{x})) \right), \quad (17)$$



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 $\text{Var}_\ell[\cdot]$ denote the mean and variance over $\ell = 1, \dots, L$ at fixed s , and let $\mathbb{E}_s[\cdot]$ and $\text{Var}_s[\cdot]$ denote the mean and variance over $s = 1, \dots, S$. We compute

$$\widehat{\text{Var}}_a(y | x) \approx \mathbb{E}_s \left[\text{Var}_\ell \left(y^{(s,\ell)}(x) \right) \right], \quad (18)$$

$$\widehat{\text{Var}}_e(y | x) \approx \text{Var}_s \left(\mathbb{E}_\ell \left[y^{(s,\ell)}(x) \right] \right), \quad (19)$$

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 $\widehat{\text{Var}}_a(y | x)$ and $\widehat{\text{Var}}_e(y | x)$, respectively.

220 3.4 Loss Function

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

$$\mathcal{L}_{\text{ELBO}} = -\mathbb{E}_{\theta \sim q_\phi(\theta)} [\log p(\tilde{y} | \tilde{x}, \theta)] + \text{KL}[q_\phi(\theta) \| p(\theta)]. \quad (20)$$

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 the KL (Kullback–Leibler divergence) term applies only to the Bayesian layers, regularizing the approximate posterior toward the prior.

For a given training pair (\tilde{x}, \tilde{y}) and a weight 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.

$$p(\tilde{y} | \tilde{x}, \theta) = \mathcal{N}(\tilde{y}; \tilde{\mu}_\theta(\tilde{x}), \text{diag}(\tilde{\sigma}_{a,\theta}(\tilde{x})^{\odot 2})), \quad (21)$$

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 $\sigma_{\text{warm}} \approx 0$. 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 DDA 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 Table C lists the hyperparameter search ranges, optimal values, and stopping criteria. Additional training diagnostics and test-set performance plots are provided in Appendix D.

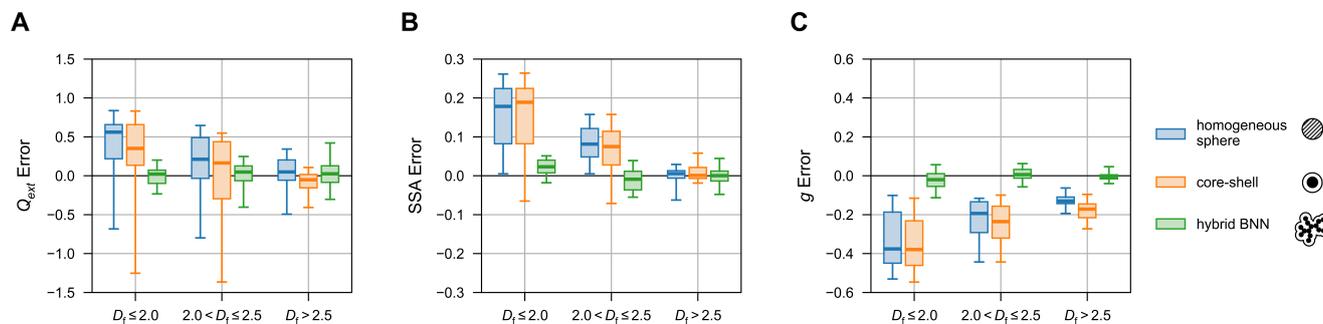


Figure 3. Accuracy of hybrid BNN compared with homogeneous sphere and core-shell approximations for calculating (A) extinction efficiency (Q_{ext}), (B) single-scattering albedo (SSA), and (C) asymmetry parameter (g). The accuracy of each model is quantified for particles with varying fractal dimension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

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

The optical properties derived from the hybrid BNN model were compared with the DDA-calculated value for each data point in the test set to evaluate the accuracy of the hybrid BNN model. We also compared the DDA-calculated optical properties with those from homogeneous sphere Mie Theory and core-shell Mie Theory, which are commonly used in large-scale models. Homogeneous sphere and core-shell Mie Theory treat BC-containing particles as homogeneous or concentric volume-equivalent spheres. We 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 $D_{f,c}$ because it accounts for the shape of the BC core and coating, as opposed to just the BC core. We separated particles into three groups based on their D_f . The first have $D_f \leq 2.0$ and represent uncoated and coated particles that have not undergone significant compaction. The second ensemble includes particles with $2.0 < D_f \leq 2.5$, and represent particles that have compacted, but their surface morphologies are not spherical. The final group have $D_f > 2.5$, and represent particles with morphologies approaching spherical.

The hybrid BNN was more accurate than homogeneous sphere and core-shell approximations for particles with $D_f \leq 2.5$ (Figure 3a). However, the accuracy of homogeneous sphere and core-shell approximations for predicting Q_{ext} improved as D_f increased. The hybrid BNN predicted Q_{ext} of particles with $D_f \leq 2.0$ 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 $2.0 < D_f \leq 2.5$ (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 $D_f > 2.5$ (median errors of 0.03, 0.05, and -0.05 , respectively).



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 have 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 hybrid BNN outputs allow for calculation of scattering and absorption efficiencies of particles (Q_{abs} and Q_{sca}). The hybrid BNN has similar accuracy to homogeneous sphere and core-shell approximations when predicting Q_{abs} across particle D_f (see Figure E1). However, the accuracy of homogeneous sphere and core-shell approximations in predicting Q_{abs} arise from canceling errors in Q_{ext} and SSA. Because Q_{abs} is given by the product of Q_{ext} and $1-SSA$, overestimation of Q_{ext} and SSA cancel, leading to increased accuracy when predicting Q_{abs} . The hybrid BNN was significantly more accurate for predicting Q_{sca} than homogeneous sphere and core-shell approximations, regardless of particle shape. More details can be found in Appendix E.

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 (d_{ve}). For these calculations, the $D_{f,c}$ 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 $D_{f,c} < 2.5$. The calculated C_{abs} of BC-containing particles was larger from the hybrid BNN than from homogeneous sphere and core-shell Mie Theory for

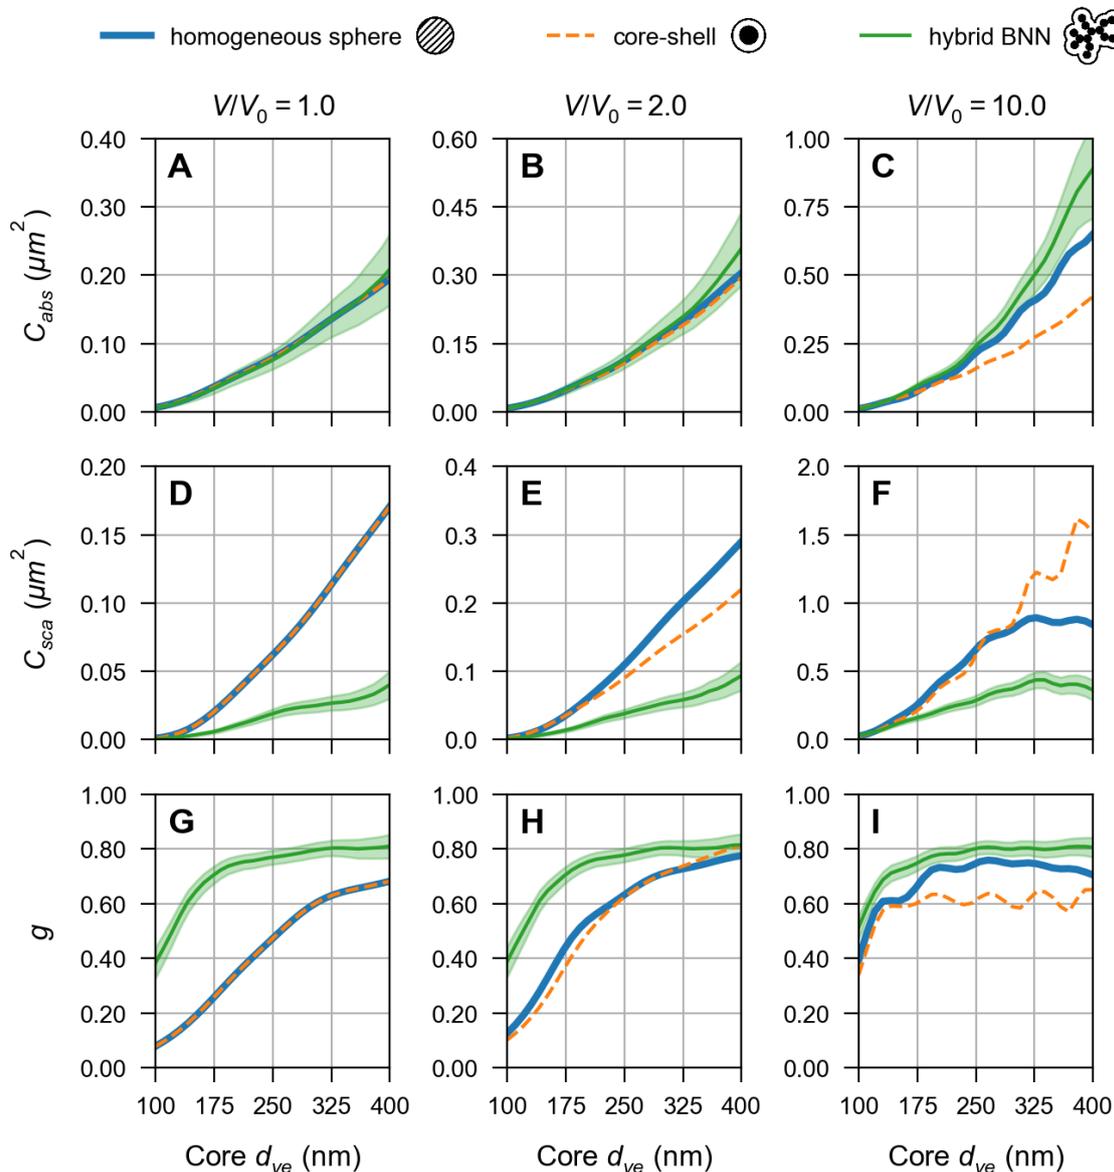


Figure 4. Absorption and scattering predictions at 532 nm using hybrid BNN compared with homogeneous sphere and core-shell approximations. (A-C) The predicted absorption cross-section (C_{abs}) of particles from the hybrid BNN is slightly larger than those using homogeneous sphere and core-shell approximations.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

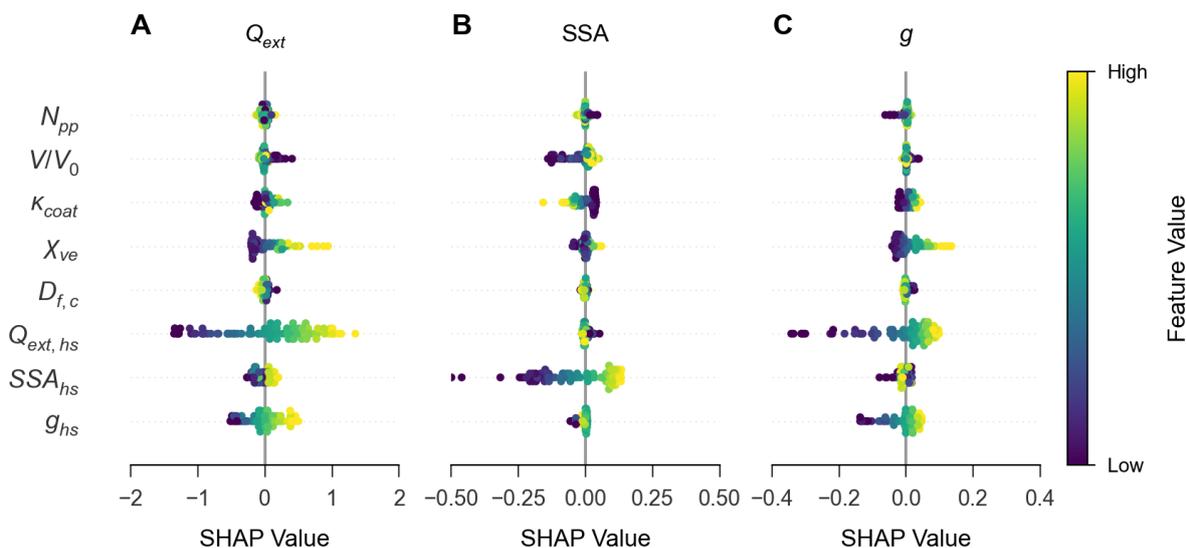


Figure 5. Shapley additive explanation (SHAP) values for extinction efficiency (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) in hybrid BNN. (A) Calculated Q_{ext} is most sensitive to χ_{ve} and $Q_{ext,hs}$ (broad SHAP distributions). (B) Calculated SSA is most sensitive to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

thickly coated particles (Figure 4a-c). This trend agrees well with previous work which showed that the absorption efficiency
 295 of large, compact particles is smaller than for large, 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 and coating amounts (Figure 4d-f), which aligns well with error trends
 observed in the training dataset. The observed lower C_{sca} for fractal particles is likely due to the optical effects arising from
 the porosity of fractal BC particles. The C_{sca} of a non-spherical particle is proportional to the effective refractive index (m_{eff})
 300 of a particle, with higher values of $|m_{eff} - 1|$ leading to higher C_{sca} . The value of $|m_{eff} - 1|$ is in turn inversely proportional
 to the porosity of a particle. 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.

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

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



310 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 value and SHAP value indicates 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; 315 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 320 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 325 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 330 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} .

335 The optical properties of an equivalent homogeneous sphere ($Q_{ext,hs}$, SSA_{hs} , and g_{hs}) also had a significant impact on all outputs of the hybrid BNN. For instance, a large input value for $Q_{ext,hs}$ led to a larger predicted Q_{ext} value. These findings show that using the optical properties of a homogeneous sphere as an input was extremely 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 340 small number of data points. The importance of $Q_{ext,hs}$, 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.

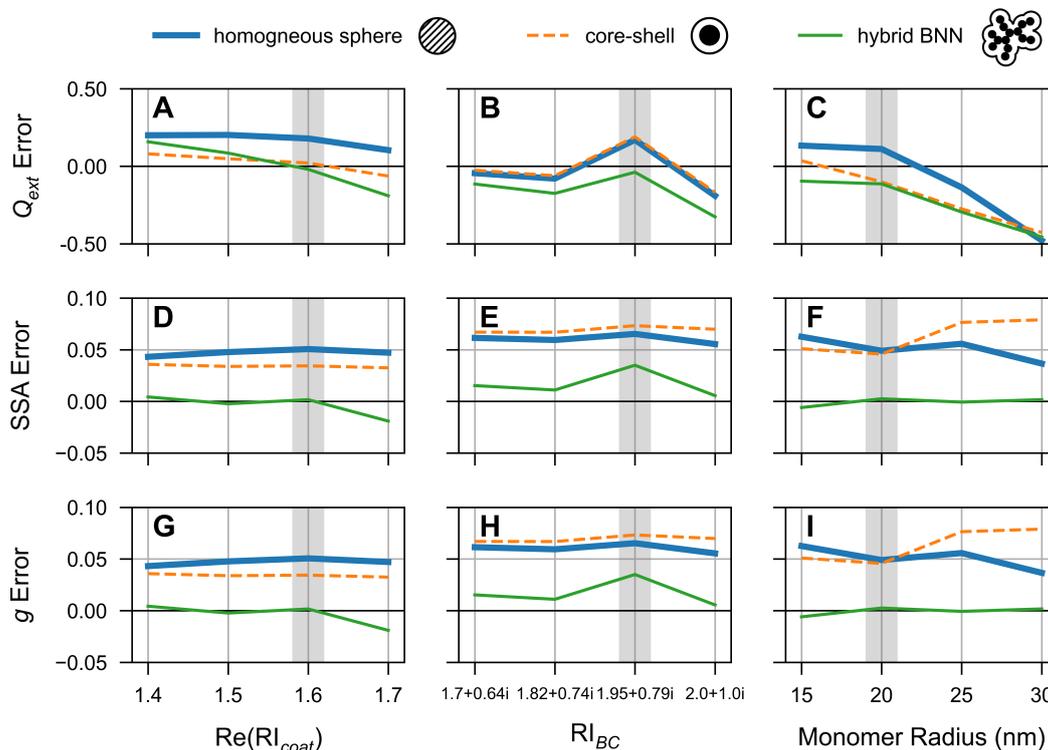


Figure 6. Mean error in extinction efficiency (Q_{ext}), single-scattering albedo (SSA), and asymmetry parameter (g) calculated by the hybrid BNN and using homogeneous sphere and core-shell approximations for a perturbed parameter ensemble. Here, error is given by the difference between predicted and true values. **(A-C)** The hybrid BNN loses accuracy when the real part of the coating refractive index ($Re(RI_{\text{coat}})$), BC refractive index (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 $Re(RI_{\text{coat}})$, RI_{BC} , and monomer radius. Shaded areas in each panel indicate values that were used in the training set and lines show average values.

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 ($Re(RI_{\text{coat}})$) of 1.6, and BC with refractive index (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, $Re(RI_{\text{coat}})$, and RI_{BC} outside of these values. We randomly selected 10 points from the evaluation dataset and systematically changed the input monomer radii in DDA simulations. The same process was repeated for simulations in which $Re(RI_{\text{coat}})$ and RI_{BC} were altered, resulting in a total of 90 additional DDA simulations. We then used the BNN to predict the optical properties of the particles and compared the BNN predictions to the DDA simulations. We benchmarked model error across $Re(RI_{\text{coat}})$, RI_{BC} ,



and monomer radii using the average difference between calculated and true values, where true values are given by DDA simulations.

Deviations in $Re(RI_{\text{coat}})$, RI_{BC} , or monomer radius from training values led to decreased accuracy by the hybrid BNN for predicting Q_{ext} , but the hybrid BNN still had comparable errors to homogeneous sphere and core-shell approximations (Figure 365 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 $Re(RI_{\text{coat}})$. 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 $Re(RI_{\text{coat}})$ 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 $Re(RI_{\text{coat}})$, RI_{BC} , and monomer 360 radii (Figure 6d-i).

It should be noted that the monomer radius, $Re(RI_{\text{coat}})$, and 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 ($Q_{\text{ext,hs}}$, SSA_{hs} , and g_{hs}), both of which are used as inputs to the hybrid BNN. Including $Q_{\text{ext,hs}}$, SSA_{hs} , and g_{hs} as hybrid BNN inputs allows for extrapolation of BNN predictions to monomer sizes and 365 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. Because epistemic uncertainty decreases as more training data is incorporated, patterns in epistemic uncertainty can be used to identify regions of the parameter space that would most benefit from additional training data. Therefore, uncertainty 370 estimates from the hybrid BNN can be leveraged to guide the design of future DDA simulations that would most effectively improve model performance.

We used the hybrid BNN to predict 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 epistemic error in optical properties as a percentage of the predicted mean value. We found that epistemic Q_{ext} uncertainty at 405 nm is highest for large, thickly coated particles with 375 $N_{\text{pp}} > 800$ and $V/V_0 > 40$ (Figure 7a). Epistemic Q_{ext} at 532 nm uncertainty is highest for thinly coated particles ($V/V_0 < 5$) and $N_{\text{pp}} < 20$ or $N_{\text{pp}} > 700$ (Figure 7b). At 880 nm, epistemic Q_{ext} uncertainty is highest for small, thinly coated particles with $N_{\text{pp}} < 50$ and $V/V_0 < 10$ (Figure 7c). For SSA and g , epistemic uncertainty across wavelengths is largest for small, thinly coated particles with $N_{\text{pp}} < 100$ and $V/V_0 < 10$ (Figure 7d-i).

Given that DDA simulations are extremely computationally expensive, this analysis can be used to identify regions of the 380 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 DDA simulations can be designed more efficiently, maximizing model improvement while reducing computational burdens.

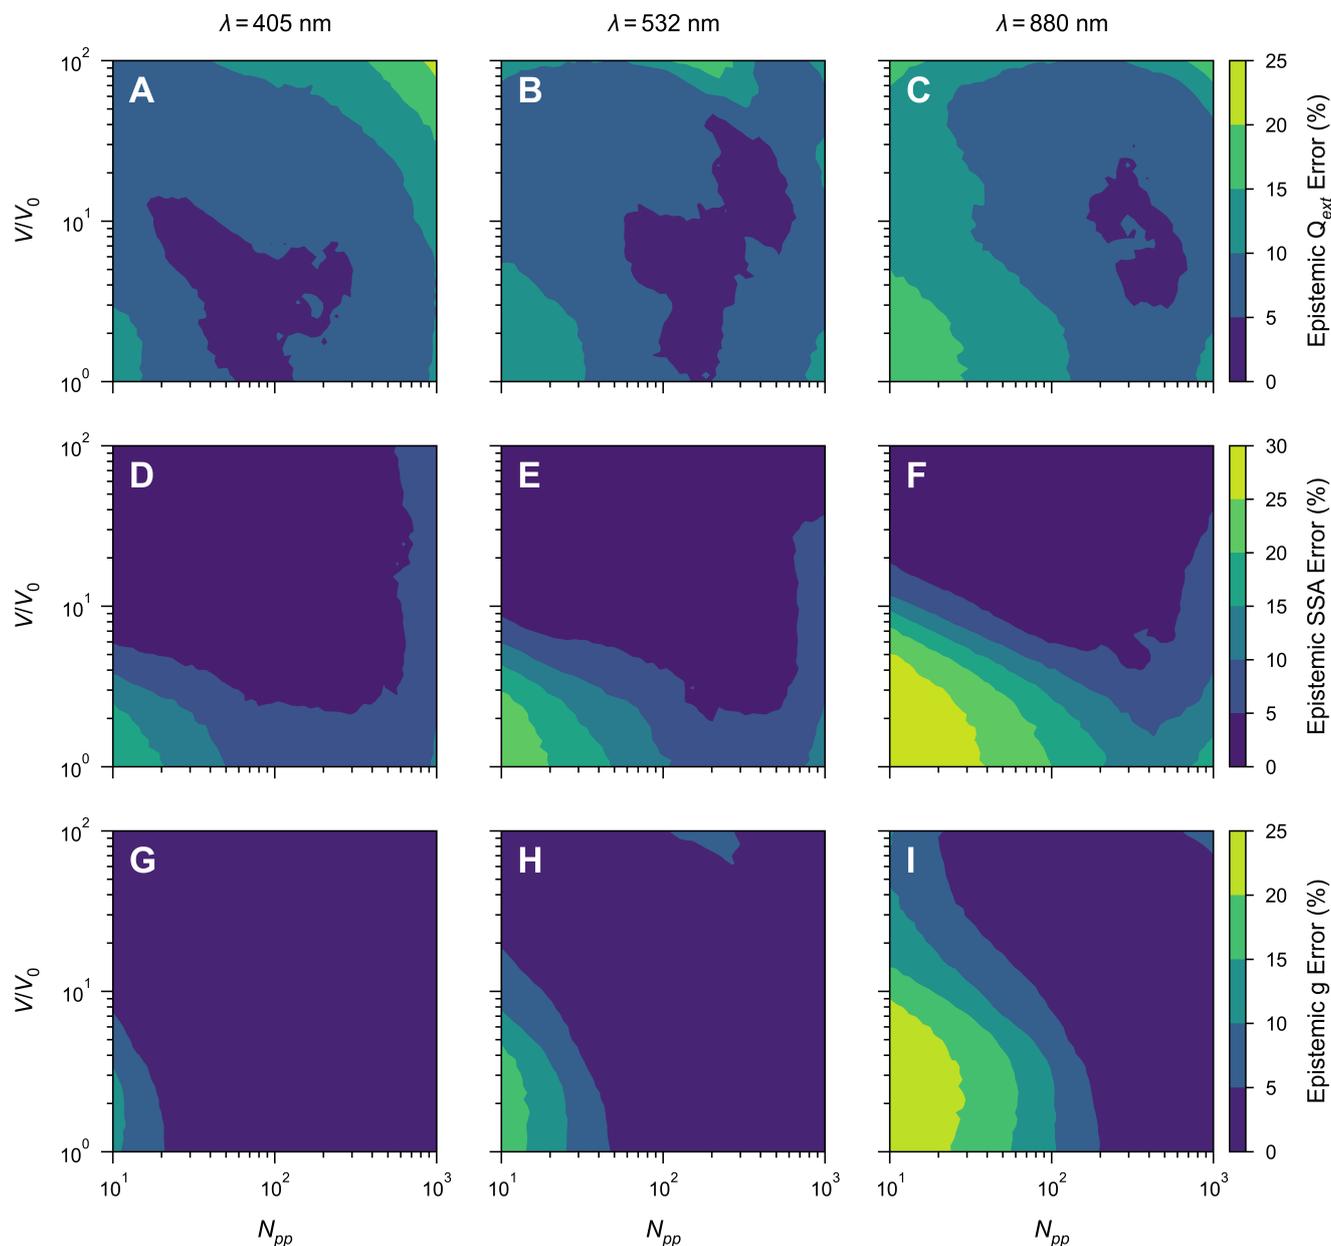


Figure 7. Epistemic uncertainty in predicted optical properties can be used for targeted future experiments aimed at reducing model uncertainty. (A-C) Epistemic uncertainty in predicted Q_{ext} is highest for large, thickly coated particles ($N_{\text{pp}} > 800$ and $V/V_0 > 40$) at 405 nm. Epistemic 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 $N_{\text{pp}} < 100$ and $V/V_0 < 10$.



9 Discussion and Conclusions

385 Efficiently predicting the optical properties of BC-containing particles remains a longstanding challenge in atmospheric mod-
els 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 bur-
dens (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 proper-
ties of BC-containing particles using machine learning-based approaches (Qin et al., 2025; Romshoo et al., 2024; Luo et al.,
390 2018). However, these studies use single-particle optics models which do not accurately capture near-field interactions be-
tween 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.
395 The model was trained using detailed single-particle optical properties from a discrete dipole approximation model, which
accurately represents the morphology and near-field interactions between monomers in BC-containing particles (Yurkin and
Hoekstra, 2011). 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
400 be calculated with low computational cost and informs 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 pre-
dicts 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
405 sphere were the most influential input features for model predictions. We also showed that inclusion of the homogeneous sphere
optical properties as an input feature allows the hybrid BNN 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. 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
410 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
415 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.

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 DDA 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 availability. The current version of DDA-BNN is available from the project website at <https://github.com/pnnl/DDA-BNN> under the MIT License. The exact version of the model used to produce the results in this paper is archived on Zenodo under DOI <https://doi.org/10.5281/zenodo.19078354> (Beeler et al., 2026).

Appendix A: 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

$$Z(x) = \frac{x - \mu_x}{\sigma_x}, \quad (\text{A1})$$

for input variables, and

$$Z(y_{tf}) = \frac{y_{tf} - \mu_{y_{tf}}}{\sigma_{y_{tf}}}, \quad (\text{A2})$$

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 B: 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 B1 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 B1 shows reliability diagrams comparing expected versus observed coverage; curves closer to the 1:1 line indicate better-calibrated predictive uncertainty.

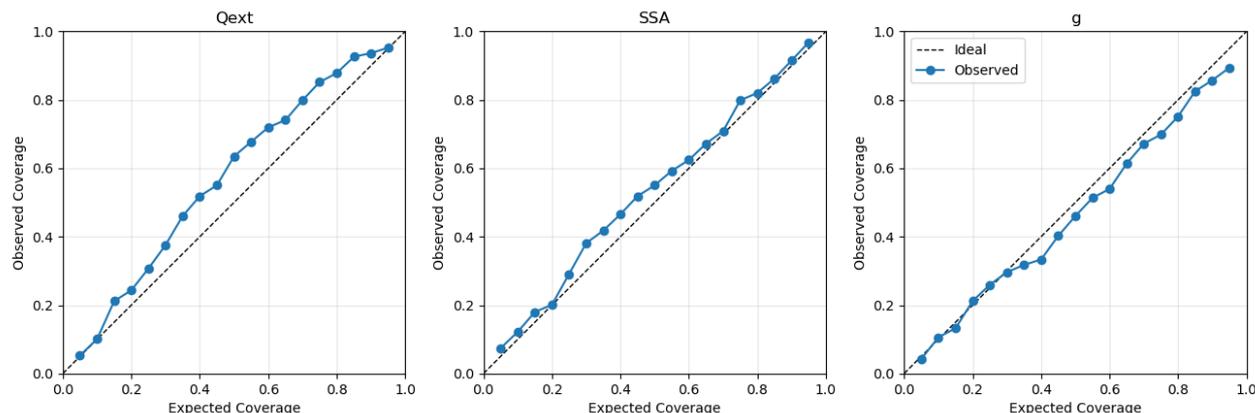


Figure B1. Reliability diagrams for hybrid BNN predictive uncertainty for Q_{ext} , SSA, and g . The dashed line indicates ideal calibration (observed coverage equals expected coverage), and the points/solid line show observed coverage. Deviations from the 1:1 line indicate miscalibration.

Table B1. Uncertainty scaling calibration results. All $\tau < 1$ indicates the raw BNN overestimated uncertainty.

Target	τ	Validation		Test	
		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%

445 Appendix C: Hybrid BNN Hyperparameter Selection

Table C1 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 $\Delta_{\text{min}} = 0.001$ 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 D.



Table C1. 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 D: Hybrid BNN Training and Test-Set Performance

Figure D1 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 D2 compares test-set predictions from the hybrid BNN (points with 90% predictive confidence intervals) and the homogeneous-sphere approximation (orange markers) against DDA reference values (dashed 1:1 line). Points closer to the 1:1 line indicate higher accuracy relative to DDA, and the BNN error bars summarize predictive uncertainty for individual test samples.

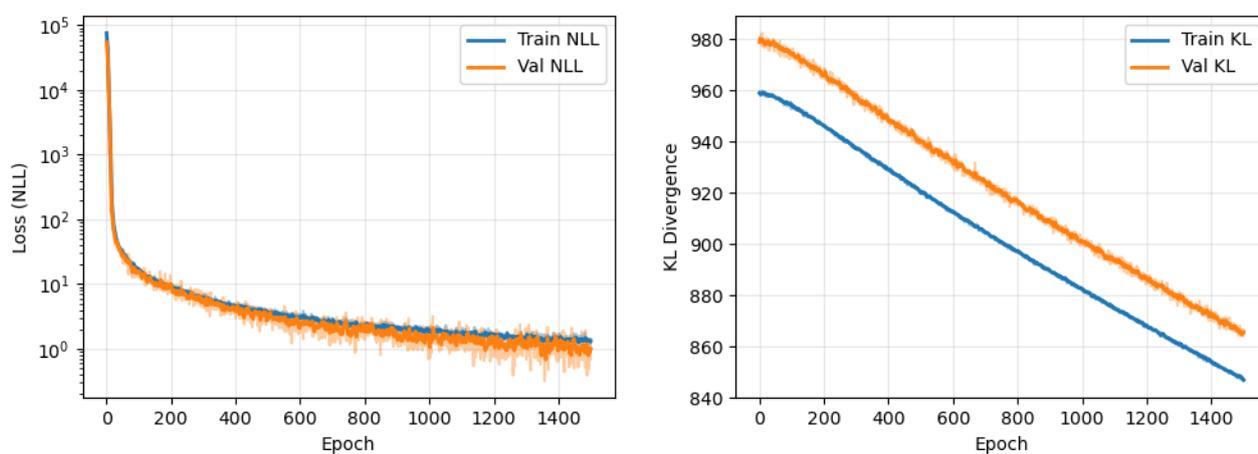


Figure D1. Training history for the hybrid BNN loss components. Left: negative log-likelihood (NLL) for training and validation sets. Right: KL divergence term for training and validation sets.

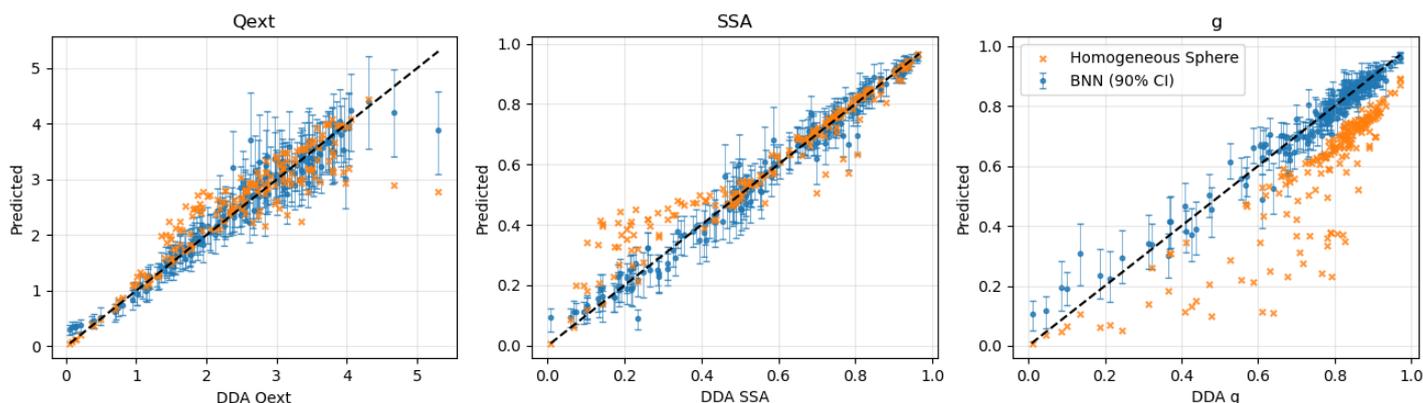


Figure D2. Test-set comparison of predicted versus DDA reference values for Q_{ext} , SSA, and g . Orange markers show the homogeneous-sphere approximation. Blue points show hybrid BNN predictions with 90% predictive confidence intervals. The dashed line indicates the 1:1 relationship with DDA.

465 Appendix E: 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 E1). The hybrid BNN calculated Q_{abs} with median error of 0.01, 0.03, and 0.01 for $D_f \leq 2.0$, $2.0 < D_f \leq 2.5$, and $D_f > 2.5$, respectively. The BNN was slightly more accurate than homogeneous sphere and core-shell approximations for $D_f \leq 2.0$ and $D_f > 2.5$, but was less accurate for $2.0 < D_f \leq 2.5$. 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 $D_f \leq 2.5$.

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.

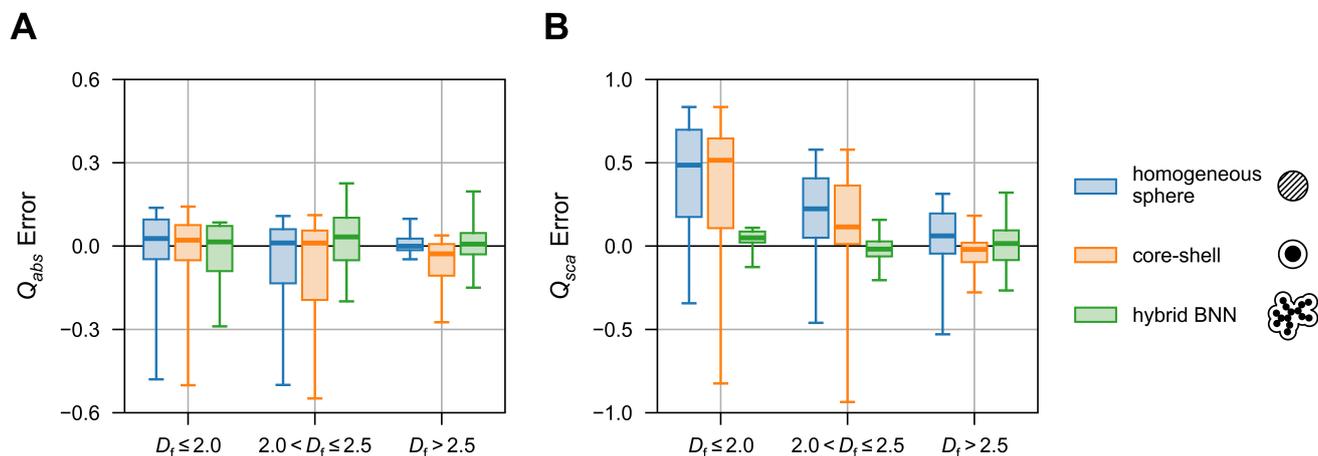


Figure E1. Accuracy of hybrid BNN, homogeneous sphere approximation, and core-shell approximation for calculating absorption efficiency (Q_{abs}) and scattering efficiency (Q_{sca}) as a function of particle shape (quantified by D_f). **(A)** The hybrid BNN is less accurate than homogeneous sphere and core-shell approximations when the shape of particles approaches that of a sphere ($D_f > 2.5$), but the hybrid BNN is more accurate for fractal particles ($D_f \leq 2.5$). **(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.

480 **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.



485 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.
- 490 Beeler, P., Donald, S., and Fierce, L.: DDA-BNN v1.0, <https://doi.org/10.5281/zenodo.19078354>, 2026.
- 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*,
495 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.:
500 Bringing uncertainty quantification to the extreme-edge with memristor-based Bayesian neural networks, *Nature Communications*, 14, 7530, 2023.
- 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 Tech-*
505 *nology*, 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.
- 510 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.
- 515 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.
- 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.
- 520 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.
- 525 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.
- Laskin, A., Laskin, J., and Nizkorodov, S. A.: Chemistry of atmospheric brown carbon, *Chemical reviews*, 115, 4335–4382, 2015.
- 530 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, 2017.
- 535 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.
- 540 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.
- 545 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.
- 550 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., Fuglestedt, J. S., Osprey, S., and Schleussner, C.-F.: Climate impacts from a removal of anthropogenic aerosol emissions, *Geophysical Research Letters*, 45, 1020–1029, 2018.
- 555 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.
- 560 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.