Response to the comments of Referee #1

We thank the anonymous referee for these important comments, and we address each point below.

Comment #1: My main concern is the treatment of particle orientation in the proposed "individual Mueller matrix selection" method. Rather than using a Mueller matrix averaged over particle orientations, the model randomly selects a particular particle realization and orientation for each scattering event and applies the corresponding Mueller matrix. The authors argue that this treatment is more representative of the physical interaction between a photon and an individual particle.

Based on my understanding, if the individual-event method and the conventional pre-averaged method represent exactly the same random particle population, I would expect the two methods to converge toward the same ensemble-mean result when the orientation-dependent extinction, scattering, and phase-matrix properties are weighted consistently. However, the manuscript reports that the two approaches give systematic differences, which can reach to even 10% as shown in Appendix A. The difference might represent genuine physics, but it might instead arise because the authors are not weighting the orientations in exactly the same way as the conventional orientation average.

In particular, the authors should demonstrate that the orientation-sampling procedure gives the correct statistical weighting to particles with different orientation-dependent optical properties. Otherwise, the difference you see compared with the standard orientation-averaged method may just come from giving different particle orientations the wrong statistical weights.

Response: We agree with the referee that generally speaking, if the two methods are intended to represent the same random particle population, then orientation selection and a rigorously implemented pre-averaged treatment should converge to the same ensemble-mean result. What we should have presented more clearly are the following three points. (1) We are interested in a general framework that can in principle handle cases in which the distribution of particle shapes (in the present work, aggregates) is not necessarily mirror symmetric. Individual nonspherical aggregates can possess shape chirality, and the population need not contain mirror-related aggregate forms in exactly equal statistical abundance. Even if the underlying population is mirror symmetric in expectation, a finite sparse realization of an optically thin layer need not contain equal numbers of mirror-related forms. In such cases, pre-averaging can still be used, but it must be performed at the level of the full dimensional vector operators. A dimensional scattering matrix averaged over realization and orientation must be constructed, and the corresponding full extinction matrix must likewise be averaged over realization and orientation rather than replaced by a scalar mean extinction cross section. These averaged matrices can retain elements that vanish for a perfectly mirror-symmetric ensemble. In this general case, the integrated scattering cross section can also depend on the incident polarization state. A rigorous pre-averaged Monte Carlo treatment must therefore allow the full Stokes vector to influence both scattering and propagation between scattering events. Our explicit-selection method, on the other hand, naturally retains the mutually correlated extinction, scattering strength, scattering matrix, handedness, and polarization response for each selected realization and orientation. (2) Our selection method additionally retains information about the variances and distributions of the Stokes-vector components

and polarization states across photon histories, as well as correlations between particle handedness/orientation and the resulting polarization. These trajectory-level statistics are removed when the particle ensemble is pre-averaged. (3) In still more general situations (for example, when the orientations of particles encountered at successive interactions are correlated), the conventional independently pre-averaged radiative-transfer description would no longer be sufficient. Our explicit-selection method, on the other hand, can be extended naturally to such cases by drawing each encountered orientation from the appropriate spatially varying and conditional or joint orientation distribution, rather than treating successive orientations as independent draws from a uniform distribution.

In the revised manuscript, we will remove the phrasing "more representative" and instead state these three more specific and more accurately phrased points clearly. In addition, we will remove the results currently shown in Appendix A (in which the averaging requirements stated in point 1 above were not fulfilled in the pre-average and additionally the weighting of the orientation selection was not consistent, though proper and consistent weighting was employed in all of the other simulations with orientation selection in the manuscript) and replace them with a demonstration of point 2, namely how our method retains information on the distribution of polarization states across photon histories. In addition, we address the referee's specific additional points below.

The manuscript states that one of 102 orientations is selected uniformly, followed by one of 360 azimuthal rotations. How are the 102 orientations distributed in three-dimensional orientation space, and do they provide an unbiased representation of randomly oriented particles?

Response: Based on Yurkin and Hoekstra (2011, Section 4.5), we specify the orientation of the scatterer relative to the laboratory reference frame by three Euler angles (α, β, and γ). The convention we use is zyz, in which α represents rotation about the z-axis, β represents rotation about the y-axis, and γ represents rotation about the new z-axis. As Yurkin and Hoekstra (2011) also point out, rotating in α is "equivalent to rotating the scattering plane without changing the orientation of the scatterer relative to the incident radiation", and thus, we set α to $0°$. Then we choose 12 values of β such that $\cos\beta$ is spaced uniformly between -1 and 1. For $\beta \neq 0°, 180°$, we choose 10 values of γ spaced uniformly between $0°$ and $360°$. For $\beta = 0°, 180°$, we include only $\gamma = 0°$. In that way, we obtain 102 combinations of β and γ (properly spaced over $\cos\beta$ and without double-counting degenerate values, i.e., unbiased) $\times$ 360 equally spaced azimuthal orientations. We did describe this in the manuscript (section 2.1), but in response to the referee's question, we will revise this description to make sure that all of these details are clear.

How is the probability of selecting a particle orientation determined? Is every orientation equally probable, or should orientations be weighted according to orientation-dependent extinction/scattering cross section?

Response: For the cases examined this study, we consider the physical orientation distribution to be uniform over rotations, i.e., we consider each physical orientation to be equally probable, similar to Mishchenko's et al. (2000) formal definition of randomly oriented nonspherical particles. It is true that the probability of *interaction of a photon* with a given orientation is also proportional to that orientation's extinction cross section relative to the ensemble-averaged extinction cross section. However, we found that, unlike the

differences among the orientation-dependent Mueller matrices and among the orientation-dependent *scattering* cross sections, the differences among the orientation-dependent *extinction* cross sections are negligible (with relative differences < 0.03), and thus we take the probability of interaction of a photon with a given orientation to also be uniform over rotations. We will clarify this very important point in the revised manuscript.

Why is the phase function for each selected orientation normalized separately? Please show mathematically that this normalization preserves the correct ensemble scattering probability.

Response: The normalization of the phase function for the chosen orientation in a given scattering event gives the correct directional probability for that scattering event. With conventional orientation averaging, the phase function would be given by:

$$P_{\text{ensemble;conventional average}}(\theta_{\text{sca}}, \phi_{\text{sca}}) = \frac{\sum_i \sigma_{\text{sca},i} P_i(\theta_{\text{sca}}, \phi_{\text{sca}})}{\sum_i \sigma_{\text{sca},i}},$$

where i is the index of the orientation, $\sigma_{\text{sca},i}$ is the scattering cross section for orientation i, and $P_i(\theta_{\text{sca}}, \phi_{\text{sca}})$ is the normalized phase function for orientation i. Over repeated events, our scheme would effectively give:

$$P_{\text{ensemble; this study}}(\theta_{\text{sca}}, \phi_{\text{sca}}) = \frac{\sum_i p_i \varpi_i P_i(\theta_{\text{sca}}, \phi_{\text{sca}})}{\sum_i p_i \varpi_i},$$

where $p_i = \frac{1}{N_{\text{orientations}}}$ is the probability of orientation i, $\varpi_i = \frac{\sigma_{\text{sca},i}}{\sigma_{\text{ext},i}}$ is the single scattering albedo associated with orientation i, and $P_i(\theta_{\text{sca}}, \phi_{\text{sca}})$ is again the normalized phase function for orientation i. Given that $\sigma_{\text{ext},i} \approx \sigma_{\text{ext}}$:

$$P_{\text{ensemble; this study}}(\theta_{\text{sca}}, \phi_{\text{sca}}) \approx \frac{\frac{1}{N_{\text{orientations}}} \cdot \frac{1}{\sigma_{\text{ext}}} \sum_i \sigma_{\text{sca},i} P_i(\theta_{\text{sca}}, \phi_{\text{sca}})}{\frac{1}{N_{\text{orientations}}} \cdot \frac{1}{\sigma_{\text{ext}}} \sum_i \sigma_{\text{sca},i}} = \frac{\sum_i \sigma_{\text{sca},i} P_i(\theta_{\text{sca}}, \phi_{\text{sca}})}{\sum_i \sigma_{\text{sca},i}},$$

which is identical to the conventional orientation-averaged phase function expression.

What should happen theoretically if both methods describe exactly the same randomly oriented particle population? If the two methods should be equivalent, why does Appendix A show a difference? If they should not be equivalent, what physical assumption makes them different?

Response: Please refer to our detailed response to the referee's general comment at the top of this document. As we stated there, in the revised manuscript, we will remove the phrasing "more representative" and instead state the three more specific and more accurately phrased points written there clearly. In addition, we will remove the results currently shown in Appendix A (in which the averaging requirements stated in point 1 above were not fulfilled in the pre-average and additionally the weighting of the orientation selection was not consistent, though proper and consistent weighting was employed in all of the other simulations with orientation selection in the manuscript) and replace them with a demonstration of how our method retains information on the distribution of polarization states across photon histories.

Comment #2: The main novelty of the manuscript concerns multiple scattering, but the comparison between the new individual-selection method and the conventional orientation-averaged method in Appendix A is limited to single scattering and one realization of the most extended aggregate. This makes it difficult to evaluate how much the proposed treatment actually affects the multiple-scattering results presented in the main text. I suggest repeating a representative subset of the full Monte Carlo calculations using both methods. For example, comparisons for one compact and one extended aggregate at an optical thickness of approximately 1, considering unpolarized, linearly polarized, and circularly polarized incident light, would be useful. Comparing quantities such as transmittance, reflectance, and selected polarization distributions would provide a much clearer demonstration of the practical impact of the proposed method.

Response: Given that we agree with the referee that if the two methods are intended to represent the same random particle population, then orientation selection and a rigorously implemented pre-averaged treatment should converge to the same ensemble-mean result (again, refer to our detailed response to the referee's general comment at the top of this document), conducting additional comparison (with appropriate weighting) over additional cases would not demonstrate the novelty of our method. Instead, as stated above, we will remove the results currently shown in Appendix A and replace them with a demonstration of how our method retains information on the distribution of polarization states across photon histories.

In addition, we will emphasize better that the results themselves with the explicit anisotropic shapes and unique polarization patterns over the various cases are also part of the novelty of the work and not solely our individual Mueller matrix selection method.

Comment #3: What is the justification of using Delta-Eddington method as a validation benchmark. I am not fully convinced that the delta-Eddington approximation provides an adequate benchmark for validating the Monte Carlo model, since delta-Eddington is itself an approximate radiative transfer treatment. The polarized-Monte-Carlo methodology cited by the authors has historically been compared against adding-doubling solutions, and such comparisons can achieve approximately percent-level agreement. Established polarized adding-doubling approaches provide a natural independent benchmark.

Response: This is an important clarification. Our comparison to delta-Eddington is *not* intended as a validation benchmark. Rather, it is intended to show how the more explicit Monte Carlo solution differs from a typical *approximate* solution often employed in large-scale (climate or global circulation) models. We mentioned this on lines 230-234 of the original submitted manuscript, but based on the referee's comment, we will emphasize it better in the revised manuscript.

Comment #4: Please carefully check the manuscript to ensure that all symbols and variables are clearly defined when they first appear and are used consistently throughout the text.

Response: As suggested, we went back over the text to check that all symbols and variables are defined and used consistently. We indeed found that the parameters $q_{n,j}$, $u_{n,j}$, and $v_{n,j}$ in

Eq. 6 were not defined explicitly in the text. In fact, we realized that they are not used anywhere after Eq. 6, and we have therefore decided to remove them. All other symbols and variables should be presented satisfactorily.

List of references that do not appear in the manuscript

Yurkin, M. A., & Hoekstra, A. G. (2011). The discrete-dipole-approximation code ADDA: Capabilities and known limitations. *Journal of Quantitative Spectroscopy and Radiative Transfer*, *112*(13), 2234-2247, https://doi.org/10.1016/j.jqsrt.2011.01.031.