

Reviewer #1 Response

AMT manuscript: *Resolving heterogeneous structure and finite scattering in clouds and precipitation with ultra-high-resolution lidar*

Authors: Grant J. Kirchhoff, Matthew Hayman, Jeffrey P. Thayer, Bryce H. Garby

The authors would like to thank the reviewer for their thoughtful consideration and feedback on the manuscript. Their input has been valuable as the resulting revisions have improved the quality of the research article.

We respond to the reviewer's comments in this document. The comments are restated below in **orange** text, our responses in **black** text, and the revised sections in the updated manuscript in **blue** and strikethrough text. In this document, the term "original lines" refers to the line numbers from the manuscript draft prior to revisions, while "updated lines" refers to the ones in the revised version.

Reviewer #1

Comments on "Resolving heterogeneous structure and finite scattering in clouds and precipitation with ultra-high-resolution lidar"

This study shows that an ultra-high-resolution lidar can reveal discrete, heterogeneous, and finite scattering features in clouds. The authors also develop a statistical framework that connects discrete-particle scattering to the classical volume backscatter lidar equation and identifies the conditions under which the volumetric interpretation is valid. They demonstrate how to estimate hydrometeor fall speed and photon flux at the single-hydrometeor scale. The topic is timely and well within scope of AMT. The ultra-high-resolution lidar observations shown in Figs. 1, 2 and 4 are striking. The manuscript is clearly written and well organized. I only have a few minor comments.

Response #1

1. Lines 265-269 and Figure 3. The parameters (e.g., particle diameter, shape parameter) are used in Figure 3 are not discussed.

The details of the parameterization are indeed not made explicit. This is the paragraph of interest (original lines 264-269):

To illustrate the implications of finite-sampling variability, we performed a simple numerical experiment using representative cloud and instrument parameters. A synthetic cloud particle population was generated from a standard droplet size distribution, i.e., gamma distribution representing size-resolved concentration: $n'(D) = N_0' \beta^\alpha \Gamma(\alpha)^{-1} D^{\alpha-1} \exp(-\beta D)$, where N_0' is the number concentration, D is the particle diameter, β and α are the respective rate and shape parameters, and $\Gamma(\alpha)$ is the gamma function. The lidar transmitter was modeled as a Gaussian beam with realistic divergence and pulse width for the system described here and range resolution of 0.1 m.

The higher concentration (500 cm^{-3}) size distribution used in the numerical experiment is shown in Fig. 1. The rate parameter $\beta = 2 \times 10^6 \text{ m}^{-1}$ and shape parameter $\alpha = 21$ were selected to be physically consistent with clouds. It is also important to point out that the purpose of the numerical experiment was to assess the impact of particle occupancy on finite-sampling variability. As a result, a physically consistent DSD was sufficient for this purpose, even if a more rigorously constrained DSD could also have been used.

The manuscript has been revised to explicitly identify these parameter values (updated lines 281-286) and further clarify the reasoning behind the experiment parameters:

Finite-sampling variability also depends on the droplet-size distribution parameterization. In this numerical experiment, the respective simulated rate and shape parameters were kept constant, i.e., $\beta = 2 \times 10^6 \text{ m}^{-1}$ and $\alpha = 21$, while the total number concentration varied from $N_0' = 500 \text{ cm}^{-3}$, i.e., warm liquid cloud, to $N_0' = 5 \text{ cm}^{-3}$, i.e., tenuous cloud or cloud edge. While it would be beneficial to test the impact of droplet-size distribution parameters on finite-sampling variability (e.g., monodisperse vs. polydisperse), we selected a single physically consistent




Figure 1: Droplet-size distribution used for the numerical experiment in the manuscript's third figure (for the 500 cm^{-3} example).

parameterization since particle occupancy is more strongly affected by concentration and range (or equivalently sample volume size).

2. Line 319: For the Gaussian mixture model, it is not clear to me how to choose the threshold/boundary of the streak. Please provide more details.

The section in question is restated here (original lines 318-322):

To estimate these slopes, a Gaussian mixture model was applied to represent each cluster, as the observed streaks are approximately elliptical in shape. The major-axis orientation of each fitted component provides an estimate of the streak's slope, from which the vertical velocity of each droplet can be inferred. The resulting velocity estimates are shown in Fig. 4d.2, where fitted ellipses and derived vertical velocities are overlaid.

The observed points in time-range space were represented probabilistically using a Gaussian mixture model with an additional uniform background component. The details in the fitting procedure and slope estimation are included in the description below, which now comprise Appendix A:

The observed points in time-range space were represented using a Gaussian mixture model with an additional uniform background component. This conceptual model is represented as

$$p(\mathbf{x}) = (1 - b) \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) + bU$$

where K is the number of Gaussian components used to represent the rain streaks, $\mathbf{x} = [t, r]^T$ contains the time and range coordinates of an observation, $\boldsymbol{\mu}_k$ is the center of the k -th Gaussian component, $\boldsymbol{\Sigma}_k$ is its full covariance matrix, π_k is its relative mixture weight, b is the fraction of observations associated with a uniform background population, and U is the probability density of that uniform background. In Fig. X, $K = 6$ represents the six rain streak clusters. The uniform component was included so that isolated background noise observations did not have to be represented by any of the Gaussian components and thereby distort the fitted Gaussian geometry.

The Gaussian centers, covariance matrices, mixture weights, and background value were optimized simultaneously by minimizing the negative mean log-likelihood:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \ln[p(\mathbf{x}_i)].$$

The slope of each streak was then obtained directly from the covariance matrix of its fitted Gaussian. Specifically, the eigenvector corresponding to the largest eigenvalue of the covariance matrix was taken as the major axis of the approximately elliptical streak. For a major-axis eigenvector $\mathbf{v}_k = [v_{t,k}, v_{r,k}]^T$, the streak slope was calculated as

$$m_k = \frac{v_{r,k}}{v_{t,k}}.$$

In addition, this sentence has been added to the manuscript to more accurately convey the meaning of the plotted ellipses in Fig. 4d.2 (updated lines 331-332):


For visualization, the fitted Gaussians were plotted as ellipses extending two standard deviations along their principal axes. Details of the fitting method are included in Appendix A.

3. Line 339 and Figure 5: please provide more details of the maximum likelihood estimation method, especially how do you estimate the flux showing in Figure 5c.2 without reading Kirchhoff et al. (2025). It might also be useful to add another curve in Figure 5c.2 to show the flux curve before the deadtime correction.

The lines in question are restated (original lines 339-341):

The data in Fig. 5c.1 were processed using maximum likelihood estimation (MLE), which enables efficient inference from sparse, clustered photon statistics while preserving time-tag resolution. The model was based on a Poisson photon-counting process modified to incorporate detector saturation effects, from Kirchhoff et al. (2025).

I added an appendix section detailing the deadtime-aware flux estimation procedure used for Fig. 5c.2. It is a useful aid to compare the deadtime-aware estimate with the standard Poisson estimate. The figure has been revised and is shown here with the updated caption:



Panel (c.2) compares the MLE range-resolved photon-flux estimate for the deadtime-aware model and Poisson model (assumes no deadtime), along with the coarse (10 m x 10 s) binned value.

This added sentence also clarifies the inclusion of the Poisson fit (updated lines 362-366):

The [deadtime-aware](#) fit explicitly accounts for saturation-induced loss of counts in the estimate, while the [Poisson fit](#) ignores deadtime effects. The [Poisson fit](#) is included for comparison to demonstrate how fitting the raw counts without accounting for deadtime results in a flux estimate that is biased low. This result is consistent with the prior demonstration of deadtime-aware photon inference at sub-meter spatial and sub-second temporal resolution in Kirchhoff et al. (2025).

The added material is now found in Appendix B, and restated below:

The fitting procedure begins with two key measured quantities: the standard photon-count histogram Y_m and the “active-fraction histogram” Z_m . Y_m represents the number of accumulated detections within the histogram bin m , while Z_m

represents the fraction of accumulated shots during which the detector was active for bin m . Here, “active” describes the state in which the detector was not disabled due to deadtime from a preceding detection event.

Next, these quantities were used to calculate the deadtime-aware Poisson negative log-likelihood:

$$\mathcal{L}_D(\lambda) = \sum_{m=1}^M (N\lambda_m Z_m \Delta t - Y_m \ln \lambda_m)$$

where λ_m is the incident photon arrival rate in bin m , N is the integrated shots, and Δt is the histogram bin size over bin number M . The ordinary Poisson likelihood is recovered when deadtime is absent, i.e., $Z_m = 1$:

$$\mathcal{L}_P(\lambda) = \sum_{m=1}^M (N\lambda_m \Delta t - Y_m \ln \lambda_m)$$

The true arrival rate, or incident photon-flux profile, λ is solved for via minimization of the loss function.

The incident photon-flux profile was represented parametrically using a Chebyshev-polynomial basis,

$$\lambda(c_J, b, t) = \exp \left[\sum_{j=0}^J c_j T_j(t) \right] + b,$$

where $T_j(t)$ is the j th Chebyshev polynomial, $c_J = \{c_j\}_{j=0}^J$ contains the fitted coefficients, J is the polynomial order, and b is a constant background photon-flux term. The exponential parameterization constrains the modeled signal contribution to remain positive.

For a specified polynomial order, the coefficients and background term were obtained by minimizing the deadtime-aware negative log-likelihood,

$$\{\tilde{c}_J, \tilde{b}\} = \arg \min_{c_J, b} \mathcal{L}_D[\lambda(c_J, b, t)].$$

Because the appropriate signal complexity was not known *a priori*, the polynomial order J was selected using holdout cross-validation. Photon detections from alternating laser shots were assigned to fitting and validation datasets, and the polynomial order producing the minimum validation loss was retained. The resulting profile $\tilde{\lambda}(t)$ is the deadtime-corrected photon-flux estimate shown in Fig. 5c.2.

The full derivation of this method and its experimental validation can be found in Kirchhoff et al. (2025).

Additionally, the relevant content in Appendix B is explicitly called out in updated lines 366-367:

Details of the fitting routine, including the deadtime-aware likelihood model, flux parameterization, and optimization procedure, are included in Appendix B.

5. Line 251: “*varation*”-> “*variation*”

The typing error has been rectified.