Review of "CrystalTrace: A Monte Carlo Raytracing Algorithm for Radiative Transfer in Cirrus Clouds with Oriented Ice Crystals," by Forster et al. 2025

Forster et al. present a new radiative transfer algorithm for cirrus clouds containing oriented ice crystals. This scheme is based on a raytracing method that can be coupled to the existing MYSTIC radiative transfer model. It handles both single and multiple scattering by oriented and randomly oriented crystals with higher angular resolution, thereby improving the simulation of sharp halos. The authors compare the method with existing data-base and pre-tabulated approaches for phase functions and radiance fields, and they show that it reproduces several optical phenomena observed during a recent field campaign.

Major comments:

1. Spectral scope and absorption
   CrystalTrace is currently limited to the visible spectral range and neglects ice absorption. Although the code successfully reproduces halo phenomena, its suitability for radiance-based applications such as fluxes and heating-rate profiles remains uncertain. Even without absorption, please compare top- and bottom-of-atmosphere fluxes with those from other models under a zero-absorption assumption to validate energy closure.
2. Crystal shapes and aggregation
   The manuscript claims that CrystalTrace supports multiple ice-crystal habits, yet only hexagonal prisms are demonstrated. Which additional shapes are implemented, what parameters (aspect ratio, bullet length, etc.) can be set, and how are aggregates treated?
3. Missing diffraction and polarization (future work)
   The authors state that diffraction and full-vector (polarized) radiative transfer will be added in a future release. Please estimate how omitting diffraction biases the asymmetry parameter (g) and top-of-atmosphere short-wave flux for thin ($\tau \approx 1$) and thick ($\tau \approx 10$) cirrus layers of different crystal sizes.

Minor comments:

1. Please indicate whether CrystalTrace can be configured for active-sensor geometries (e.g. zenith-pointing or spaceborne lidar).
2. For the benefit of prospective users, consider supplying a concise example input deck (MYSTIC / libRadtran format) that defines an ice cloud with user-selected aspect ratio and orientation parameters.
3. Include a benchmark case with only randomly oriented crystals, comparing CrystalTrace results to those from an established solver. This will demonstrate that the new code reproduces standard solutions when orientation effects are disabled.

4. Clarify whether the current implementation can handle mixed-phase clouds that contain both ice and liquid water.

Specific comments:

p.1 l-14: Please quantify the impact of ice-crystal orientation on the radiation budget (e.g. ΔSW CRE).

p.3 l-46: The code is said to build on Macke (1994). Please spell out the improvements and novel aspects relative to that work. Section 2.3 compares CrystalTrace with the bug-fixed version of GO; clarify the differences between the two algorithms.

p.3 l-47: If a broader wavelength range than 400–700 nm is simulated, will the results near 400 nm and 700 nm remain consistent with other algorithms?

p.7 l-144: Is photon termination (e.g., Russian-roulette) applied to prevent endless scattering in very thick ice clouds?

p.8 l-173: The sentence contains two instances of “with”; please remove one.

p.8 l-184: The authors show that CrystalTrace is unaffected by scattering-angle resolution. What is the computational-cost difference between CrystalTrace and the pre-computed look-up table?

p.9 Figure 6: The legends list “4.0e-04” and “2.1e-03”, but the meaning of these values is not explained.

p.10 Figure 7c: The background (non-22°) radiance is slightly higher in CrystalTrace. Was ice absorption enabled in the MYSTIC-YANG control run? If not, what explains the difference?

p.10 l-200: Please provide the standard deviation of radiances for MYSTIC-CrystalTrace versus MYSTIC-YANG over several independent runs.

p.10 l-202: Runtime is reported for only 1×10^5 photons—insufficient for low noise in a hemispheric simulation. Please also give the runtime (and peak memory) for an equivalent MYSTIC-YANG run to substantiate the claimed efficiency.

p.13 l-233: How many photons were used to obtain adequate aerosol-scattering statistics?