Peer Review: “LACRIT: an automated lidar framework for detecting cirrus and contrails and retrieving their radiative properties”

General assessment

The manuscript presents a well-engineered, carefully documented layer-detection and retrieval framework (LACRIT) combining transmittance-based COD estimation with a modified particular-integration (PI) method for LR retrieval, validated through internal consistency checks (COD_{trans} vs. COD_{PI} agreement within 1–4%). The instrumental description and processing chain (SNIP baseline, MAD-based adaptive thresholding, hysteresis detection, regularized LR inversion) are sound and clearly presented. However, I have a significant methodological concern regarding the interpretation of the COD–VLDR relationship (Sect. 4.0.1, Fig. 6) that should be addressed before publication, as it affects one of the paper’s central physical claims.

We interpret the authors’ VLDR as the ratio between the depolarized (cross-polarized) lidar return and the polarized (co-polarized) lidar return, suitably calibrated to a molecular value in a region of the profile assumed to be free of aerosol. If this is indeed the definition adopted, it would be helpful for the authors to state it explicitly in Sect. 2.1 or 3.1, together with the corresponding equation, since this definition is central to the concerns raised below. From now on, we will use the more widely used notation δ_v to indicate VLDR.

Major comment: the reported COD–δ_v dependence may be largely a definitional/mixing artifact, not (only) a microphysical signal

The authors report a systematic, sublinear (power-law, saturating) relationship between COD and δ_v at 532 nm ($a = 0.56$, $b = 0.157$, $R^2 = 0.37$), which they interpret in terms of genuine microphysical processes — ice-crystal habit, particle size distribution, preferred orientation, and multiple scattering — with sensitivity of depolarization to COD decreasing as clouds become optically thicker. I believe this interpretation needs to be qualified, because **the volume linear depolarization ratio δ_v is, by construction, not a single-particle property** — unlike the particle depolarization ratio δ_p discussed as istance by Cairo et al. (1999) — but a backscatter-weighted mixture of the molecular and particulate contributions within the probed volume. This mixture necessarily depends on the *amount* of particulate matter present (i.e., on backscatter/extinction, and hence indirectly on COD), independently of any change in the intrinsic scattering properties of the ice particles themselves.

The mixing formula

For a two-component medium (molecules + particles) sharing the same range bin, the parallel and cross-polarized signals are additive, hence:

$$\delta_v(r) = \frac{\delta_{\mathrm{mol}}\, \beta_{\mathrm{mol}}(r) + \delta_p(r)\, \beta_{\mathrm{p}}(r)}{\beta_{\mathrm{mol}}(r) + \beta_{\mathrm{p}}(r)} \tag{1}$$

where

$$\beta_{\mathrm{mol}} = \beta_{\mathrm{mol}}^{\mathrm{cross}} + \beta_{\mathrm{mol}}^{\mathrm{par}}, \qquad \beta_p = \beta_p^{\mathrm{cross}} + \beta_p^{\mathrm{par}}, \qquad \delta_p = \frac{\beta_p^{\mathrm{cross}}}{\beta_p^{\mathrm{par}}}, \qquad \delta_{\mathrm{mol}} = \frac{\beta_{\mathrm{mol}}^{\mathrm{cross}}}{\beta_{\mathrm{mol}}^{\mathrm{par}}} \tag{2}$$

Introducing the backscatter ratio $R(r) = \dfrac{\beta_{\mathrm{mol}}(r) + \beta_{\mathrm{p}}(r)}{\beta_{\mathrm{mol}}(r)}$, this can be rewritten as

$$\delta_v(r) = \frac{\delta_{\mathrm{mol}}}{R(r)} + \delta_p(r)\left(1 - \frac{1}{R(r)}\right) \tag{3}$$

This expression shows that, **even for a constant, altitude- and time-independent particle depolarization** δ_p, δ_v is a monotonically increasing, concave (saturating) function of R: it tends to the small molecular value δ_{mol} ($\approx$ 0.003–0.01, (Behrendt and Nakamura, 2002; Kong et al., 2023)) when particulate backscatter is negligible ($R \to 1$, subvisible/very thin layers) and asymptotically approaches the true particle value δ_p as R grows ($R \to \infty$), see as istance Adachi (2001). Since R is tightly linked to particulate extinction/backscatter, and therefore to COD (COD $= \mathrm{LR} \int \beta_{\mathrm{par}}\, dr$), a **sublinear, saturating COD–δ_v curve is a mathematical consequence of the volume-depolarization definition itself**, not necessarily evidence of ice-crystal growth, aggregation, or habit change.

This is precisely the qualitative shape reported in Fig. 6 (power law with $b < 1$, flattening at high COD), and it occurs in exactly the regime where the effect should be strongest: subvisible to thin cirrus (COD $<$ 0.05–0.3), where the molecular contribution to the volume signal is far from negligible — the very regime this paper targets and resolves.

A more quantitative check: comparing the empirical fit against the theoretical mixing curve

It is worth pushing this point further than a qualitative remark, because the authors have all the ingredients needed for a direct, almost cost-free test. Within a single, geometrically thin, roughly homogeneous layer, COD and R are approximately related by

$$\bar{\beta}_{\mathrm{par}} \cdot H \approx \frac{\mathrm{COD}}{\mathrm{LR}} \quad \Rightarrow \quad R \approx 1 + \frac{\mathrm{COD}}{\mathrm{LR} \cdot \beta_{\mathrm{mol}} \cdot H} \tag{4}$$

so that, substituting into the mixing formula above with $\delta_{\mathrm{mol}} \to 0$ for simplicity,

$$\delta_v(\mathrm{COD}) \approx \delta_p \left(1 - \frac{1}{1 + k \cdot \mathrm{COD}}\right) = \delta_p \cdot \frac{k \cdot \mathrm{COD}}{1 + k \cdot \mathrm{COD}}, \qquad k = \frac{1}{\mathrm{LR} \cdot \beta_{\mathrm{mol}} \cdot H} \tag{5}$$

This is a **saturating hyperbolic (Michaelis–Menten-type) function of COD with only two free parameters** (δ_p and k), not a three/four-parameter empirical power law. It would be a stronger and more physically motivated test to add this curve as an explicit candidate in the AICc model comparison already performed in Sect. 4.0.1 (the framework used, comparing "several nonlinear candidate models," is exactly the right tool for this — it just seems to be missing the one model that is actually derived from the definition of δ_v itself). If this physically constrained curve fits the data in Fig. 6 comparably well to (or better than) the freely-fitted power law, that would be fairly direct evidence that the observed relationship is dominated by the mixing effect rather than by independent microphysical evolution, and the AICc ranking currently reported would need to include it for the comparison to be complete.

Practical implication and suggested revision

The good news is that the authors already compute, internally, all the quantities needed to disentangle this effect, since $R(r)$ is essentially the same normalized-signal-to-molecular-baseline ratio ($S_{\mathrm{PI}}/B_{\mathrm{PI}}$, or $\tilde{S}(r)$ vs. the Rayleigh reference) already used for the $\mathrm{COD}_{\mathrm{PI}}$ retrieval (Eq. 8). I would recommend:

- **Retrieve the particle depolarization ratio** δ_p (not just δ_v) for each detected layer core, using the un-mixing formula above with R estimated from the same normalized elastic/molecular ratio already computed in the LACRIT pipeline (this is a very modest addition to the existing processing chain, not new instrumentation or new theory).
- **Re-run the COD vs. depolarization regression using δ_p instead of δ_v.** If the sublinear trend largely disappears (i.e., δ_p shows much weaker or no dependence on COD), this would demonstrate that the originally reported relationship in Fig. 6 was dominated by the trivial mixing effect rather than by microphysics, and the paper's interpretive discussion (lines 296–312) would need to be substantially revised.
- If a **residual** COD-dependence of δ_p remains after this correction, that residual is the part that can legitimately be attributed to microphysical evolution (ice-crystal growth, aggregation, habit change, or genuine multiple scattering — see below) and should be the basis for the physical discussion currently built on δ_v.
- This distinction should also be revisited for the **Layer 1 vs. Layer 2 comparison** (Table 1) and the **fresh vs. aged contrail comparison** (Table 2, VLDR ≈ 0.20 for the fresh contrail vs. 0.43 at similar-but-not-identical COD): since COD differs somewhat between the compared cases (0.06–0.08 vs. 0.12), part of the reported depolarization contrast could still reflect the R-dependence of δ_v rather than a purely microphysical fresh-vs-aged difference. Given that the paper's key narrative result (young vs. aged contrail depolarization contrast) rests on this comparison, it would strengthen the conclusions considerably to show it survives in terms of δ_p.
- A useful sanity diagnostic, which costs nothing computationally, would be to **plot δ_v directly against** R (rather than against COD) for the same dataset used in Fig. 6. Since R is the variable that actually and directly controls the mixing, if the collapse onto a single curve is *tighter* in δ_v–R space than in δ_v–COD space, that in itself would be informative evidence that mixing, not microphysics, is the leading-order driver of the pattern in Fig. 6 (COD and R are correlated but not identical, since COD also integrates the vertical layer thickness H, which varies between the layers/cases compared here from ~ 350 m to ~ 585 m). This behaviour is seen most clearly by plotting δ_v against $1 - 1/R(r)$ rather than against $R(r)$ itself: according to the mixing formula above, this transformation linearises the relationship, so that a plot of δ_v versus $1 - 1/R(r)$ should collapse onto a straight line of slope δ_p and intercept δ_{mol} whenever δ_p is approximately constant across the dataset, exactly as shown by Adachi (2001). Reproducing this linearised plot for the present dataset would offer a direct, low-effort visual test of whether the curvature seen in Fig. 6 is dominated by mixing.

Multiple scattering as a separate, compounding effect

Independently of the volume/particle mixing issue above, **multiple scattering (MS)** can also introduce a genuine (not purely definitional) increase of apparent depolarization with increasing optical depth/extinction, for lidars with finite field of view (FOV): photons undergo more depolarizing forward-scattering events before being detected as the medium becomes optically thicker, an effect well documented since Eloranta's and Hu et al.'s multiple-scattering treatments and widely discussed in the cirrus lidar literature the authors already cite (e.g., Zhou et al., 2012). This is conceptually distinct from the mixing artifact discussed above (MS is a radiative-transfer effect operating even on a single, hypothetically monodisperse particle population), but it acts in the *same direction* — increasing apparent depolarization/apparent lidar ratio behavior with optical depth — and would compound with the mixing effect rather than being distinguishable from it using COD–δ_v or COD–δ_p data alone.

I could not find in Sect. 2.1 a statement of the CE376's receiver FOV or divergence, which would let a reader estimate whether MS is expected to be significant at the COD/range regime sampled here (COD up to ≈ 0.9, ranges 8–14 km). Given the biaxial, compact-telescope configuration described, MS is plausibly modest for the thinner cases but should not be assumed negligible for the higher-COD points driving the curvature of the fit in Fig. 6. I recommend that the authors:

- state the receiver FOV/divergence explicitly, and
- provide at least an order-of-magnitude estimate (or a simple correction, e.g., following Eloranta's single-scattering-equivalent approximation) of the expected MS contribution to δ_v at the highest COD values sampled, to support the claim (lines 305–307) that gradual saturation is attributable to "multiple-scattering effects" — a claim which is currently asserted by reference to other studies rather than demonstrated from the authors' own data/instrument geometry.
- Note also that MS depends on the actual physical range to cloud base (larger footprint at 10–13 km than in near-range studies), and on particle size, and so could differ systematically between Layer 1 ($\sim$10 km) and Layer 2 ($\sim$12.5 km, colder, likely smaller/different habit ice crystals) in Table 1 — potentially confounding the LR/VLDR contrast already attributed there purely to microphysics (lines 267–277).

Consequence for the AICc model-selection exercise

Because the sublinear shape being fitted (power law, $b = 0.157$) is qualitatively exactly what the mixing formula predicts even for a *constant* δ_p, the AICc-based comparison of nonlinear models for COD–δ_v (Sect. 4.0.1) risks selecting and physically over-interpreting curvature parameters that are, at least in part, forced by the definition of volume depolarization itself, rather than freely diagnostic of cloud microphysics. I would suggest repeating this model-selection exercise on δ_p once available, or at minimum explicitly caveating that the moderate R^2 (0.37) and the specific curvature parameter b cannot be interpreted as purely microphysical without first removing the R-mixing dependence.

I would add one further, purely statistical, caveat to this exercise: the $N = 63$ points entering the fit in Fig. 6 are drawn from time series averaged over consecutive 5-minute windows within a small number of continuous observation episodes (25 Dec 2022, 6 Sep 2025), rather than from independent cloud realizations. Consecutive 5-minute retrievals of the same evolving layer are very likely autocorrelated in both COD and δ_v (a slowly thickening or thinning layer will produce a run of similar, non-independent values). This inflates the effective sample size used implicitly by AICc and by the reported R^2, in much the same way that testing many non-independent hypotheses inflates the apparent number of "significant" findings if left uncorrected. A rough estimate of the effective number of independent samples (e.g., from the integral autocorrelation time of the COD or δ_v series) would allow a more honest assessment of how much weight the moderate $R^2 = 0.37$ and the AICc ranking can actually bear.

Minor comments

- Please confirm the definition and typical instrumental value used for δ_{mol} (Rayleigh depolarization), since this directly enters the correction discussed above and also affects the low-COD asymptote visible in Fig. 6.
- It would help the reader if the manuscript stated explicitly how δ_v is computed from the raw parallel/cross-polarized channels (e.g., whether a polarization gain ratio / calibration constant has been applied, and by which method — e.g., $\pm 45°$ rotator calibration). This is standard practice but is not described in Sect. 2.1 or 3.1, and any residual cross-talk or gain-ratio calibration error would add a further, purely instrumental, source of the reported COD–δ_v curvature (since calibration residuals often scale with signal strength, hence with COD). More generally, the authors should describe in detail their depolarization calibration procedure, including the calibration method used, the frequency at which it is performed, and the assumed molecular depolarization reference value δ_{mol} since the reliability of the entire δ_v/δ_p analysis rests on this calibration.
- The CE376 lidar also acquires elastic backscatter at 808 nm (Sect. 2.1), but this second channel does not appear to be exploited anywhere in the retrieval or in the case-study analysis of Sect. 4. Did the authors consider using the co-located 532/808 nm measurements to build a backscatter-related colour index (e.g., an attenuated colour ratio or wavelength-dependent backscatter ratio)? Such an index is well known to provide additional, largely independent constraints on particle size and habit, complementary to VLDR and LR, and could help disentangle whether the COD–VLDR relationship discussed above is genuinely microphysical or dominated by the molecular/particulate mixing effect — e.g., by checking whether the colour index also varies systematically with COD/R in the same layers, or instead remains stable while VLDR changes. Even a qualitative discussion of the 808 nm channel's potential in this respect, or an explanation of why it was not used here, would strengthen the manuscript.
- Table 2 reports several entries with no stated uncertainty (e.g., L2 04:12–04:20 row); please clarify whether these are single retrievals ($N = 1$) rather than 5-min block averages, since this affects how confidently the fresh-contrail contrast can be discussed.
- Sect. 2.2 has a minor typo ("all-sky all-sky camera").

Recommendation

Minor-to-moderate revision. The instrumental and retrieval framework is solid and the internal COD consistency check is convincing. However, the physical interpretation of the COD–depolarization relationship — one of the paper's headline results — should be revisited using the particle (not volume) depolarization ratio, tested where possible against the explicit mixing-derived functional form proposed above, and the potential compounding role of multiple scattering and polarization-calibration residuals should be at least quantitatively bounded, before the microphysical interpretation currently given in Sect. 4.0.1 can be considered fully supported by the data.

References

Adachi, H.: Retrieval of particle depolarization ratio from lidar-measured volume depolarization ratio using a linearized backscatter-ratio approach, *Applied Optics*, 40, 2001.

Cairo, F., Di Donfrancesco, G., Adriani, A., Pulvirenti, L., and Fierli, F.: Comparison of various linear depolarization parameters measured by lidar, *Applied Optics*, 38, 4425–4432, https://doi.org/10.1364/AO.38.004425, 1999.

Behrendt, A. and Nakamura, T.: Calculation of the calibration constant of polarization lidar and its dependency on atmospheric temperature, *Optics Express*, 10, 805–817, https://doi.org/10.1364/OE.10.000805, 2002.

Kong, Z., Yang, X., and Mei, L.: Theoretical and experimental investigation of the molecular depolarization ratio for broadband polarization lidar techniques, *Optics Express*, 31, 24897–24913, https://doi.org/10.1364/OE.492240, 2023.