

Response to Reviewer #2

Revised Manuscript

**LACRIT: an automated lidar framework for detecting cirrus and contrails
and retrieving their radiative properties**

Authors

Andres E. Bedoya-Velásquez^{1,*}, Claire Sarrat², Céline Parzani², and Romain Ceolato¹

¹ONERA, DOTA, Université de Toulouse, 31000 Toulouse, France

²ONERA, DTIS, Université de Toulouse, 31000 Toulouse, France

*Corresponding author: andres.bedoya@onera.fr

Journal: Atmospheric Measurement Techniques
Manuscript ID: egusphere-2026-3056
Revision: R2
Date: September 9, 2026

General response

We thank the reviewer for the detailed and constructive comments and greatly appreciate the time and effort devoted to reviewing our work.

In the revised manuscript, most changes introduced in response to the reviewer’s comments are highlighted in red to facilitate their identification and allow the revisions to be easily tracked. Sections 3.2 and 3.3, however, were extensively restructured and, in several parts, almost entirely rewritten, while the Appendix was newly developed as part of the revision. Given the extent and length of these modifications, highlighting them in red would have resulted in nearly entire sections being colored, defeating the purpose of using color to identify specific changes. These sections and the Appendix have therefore been left in the standard black text, while revisions made elsewhere in the manuscript remain highlighted in red.

For completeness, and to facilitate the evaluation of the responses below, the substantially revised Sections 3.1, 3.2 and 3.3 and the newly developed Appendix are reproduced at the end of this response document. These sections contain methodological developments and validation analyses that are referred to extensively throughout our responses. They are provided here solely to allow the reviewer to assess the corresponding revisions in context, as the revised manuscript itself cannot be attached at this stage of the submission procedure.

Notation note: To ensure consistent notation throughout the revised manuscript, VLDR has been systematically replaced by δ_v . In the responses below, however, we retain the term VLDR where appropriate to remain consistent with the notation used in the original manuscript and avoid ambiguity.

1 Summary and general assessment

Reviewer comment

The authors present a new method called LACRIT for the retrieval of the cloud optical depth, lidar ratio, and volume linear depolarization ratio of cirrus clouds formed by contrails. LACRIT appears to be a synergistic retrieval using a micro-pulse lidar, an all-sky camera, and a microwave radiometer. The all-sky camera and the microwave radiometer play secondary roles. The former is used to visually verify the presence of the cirrus clouds and track their movement and the latter is used as an alternative for the meteorological parameters used for the molecular profile when ERA5 reanalysis data are not available. Two cases are presented showing an application of the LACRIT method and the results are discussed and compared with the literature.

I have focused my review on the LACRIT method as my expertise is not in cloud micro-physics. I would recommend the paper for publication after dealing with some serious issues regarding the optimal estimation techniques deployed for the LACRIT retrieval and after revision of Section 3.

Response

We thank the reviewer for the careful assessment of the manuscript and for focusing particularly on the methodological aspects of LACRIT. We agree that the original version of Section 3 did not describe the retrieval, its constraints, and its uncertainty assessment with sufficient clarity.

A synthetic validation of the retrieval had already been initiated as part of the methodological developments planned for a subsequent climatological study. However, after reconsidering the scope of the present manuscript in light of the reviewer’s comments, we concluded that such a validation is fundamental to establishing the robustness of LACRIT itself and should therefore be included in the present work rather than deferred to a future application study.

Accordingly, the methodological and validation sections have been substantially reorganized. Section 3 now provides the complete methodological definition of the LACRIT processing chain, including the cloud-detection procedure, the molecular-referenced transmittance retrieval, the modified PI lidar-ratio inversion, the uncertainty weighting, and the role of static and temporal regularization. Section 4 now begins with the principal results of an independent synthetic forward–inverse validation, in which cloud geometry, COD, and LR are prescribed independently and subsequently retrieved from simulated elastic-lidar profiles under deterministic and controlled-noise conditions.

To keep the main manuscript focused while preserving reproducibility, the more technical elements of the validation have been transferred to Appendix A – E. These include the construction of the synthetic atmospheric state and elastic forward model, the detailed Monte Carlo implementation and truth-matching metrics, the static-prior and temporal-regularization sensitivity experiments, and a dedicated discussion of the scope and limitations of the synthetic validation.

This revised organization allows cloud detection, boundary localization, COD retrieval, LR retrieval, and the influence of regularization to be assessed separately, while avoiding an overly technical expansion of the main methodology. The corresponding changes are described in detail in our responses to the specific comments below.

2 General comments

2.1 Validation of the LACRIT retrieval

Reviewer comment

The transmission method and the lidar ratio retrieval method of Cadet et al. are well established retrieval techniques for cloud properties. A big part of the novelty of this paper is the optimized retrieval of LACRIT. In its current state, the paper is lacking a proper validation of the new technique. The authors can consider applying the technique on synthetic data and compare their retrievals with the Cadet method to convince that LACRIT is indeed better or at least equally good.

Moreover, the retrieval technique and the uncertainty calculation are not adequately described in Section 3.1 (more specific comments below).

Response

We thank the reviewer for this important comment. We agree that the original version of the manuscript relied too strongly on the agreement between the transmittance and PI optical-depth estimates as evidence of retrieval performance. Because these two quantities participate in the same LR optimization, their agreement is more appropriately interpreted as an internal consistency condition than as an independent validation.

To address this point, we have substantially revised Sec. 3 and introduced a truth-based synthetic forward–inverse validation of the complete LACRIT processing chain. The experiment begins from an atmosphere for which cloud geometry, cloud optical depth (COD), and particulate lidar ratio (LR) are prescribed exactly. These quantities are used only to generate the corresponding single-scattering elastic-lidar signal. The resulting synthetic profiles are

then processed using the same normalization, cloud detection, molecular-referenced transmission, modified PI retrieval, and regularization procedures described for LACRIT. The prescribed truth is accessed only after completion of the retrieval and is therefore not used to define cloud boundaries, COD, or LR.

The deterministic validation comprises 20 independent combinations spanning $\text{COD} = 0.02\text{--}0.60$ and $\text{LR} = 20\text{--}50$ sr. Measurement-noise sensitivity is subsequently evaluated at target SNR values of 50, 20, 10, 5, and 2, using 50 independent Monte Carlo realizations for each COD–LR combination. Detection performance, cloud-boundary localization, transmission COD, and LR accuracy are evaluated separately in order to distinguish successful layer identification from the accuracy of the subsequent optical retrieval.

Under noise-free conditions, the transmission-based COD is recovered with negligible error, while the unregularized `DATA_ONLY` LR retrieval yields an overall RMSE of 5.54 sr. The deterministic experiment also reveals an increasing positive LR bias toward the largest prescribed COD values, which is now explicitly identified as a limitation of the present PI-based formulation under the assumptions of the synthetic framework. With increasing noise, truth-matched cloud detection remains 99.4% even at the lowest tested SNR of 2, whereas the uncertainty of the optical retrieval increases substantially. For example, the `DATA_ONLY` LR RMSE increases from 5.54 sr in the deterministic case to 42.29 sr at SNR 2. These results are now presented in the main Results section, while the detailed forward model, Monte Carlo implementation, truth-matching criteria, regularization experiments, and validation assumptions are documented in the Appendix.

We have also substantially revised the description of the retrieval and its uncertainty treatment in Sec. 3. The uncertainty of the molecular-referenced transmission estimate is now explicitly propagated into the LR optimization through σ_{COD} , which weights the PI–transmission mismatch according to the uncertainty of the transmission constraint. The synthetic Monte Carlo experiments then provide a truth-based assessment of how progressive signal degradation propagates through layer detection, boundary localization, COD, and LR retrieval.

Regarding the suggested comparison with the method of Cadet et al., we have clarified the scope of the present contribution. LACRIT does not aim to demonstrate that the PI formulation itself is superior to that of Cadet et al. (2005); rather, its LR retrieval is a modified implementation of that formulation embedded within an automated framework including layer detection, molecular-referenced transmission, uncertainty weighting, and optional regularization. A direct implementation-to-implementation comparison would therefore primarily assess differences in the practical realization of the PI approach, whereas the prescribed-truth experiment allows retrieval accuracy and bias to be quantified directly. We consequently use synthetic truth as the primary validation benchmark. Forward and backward Fernald–Klett inversions are nevertheless retained as complementary consistency diagnostics.

We believe that these revisions provide the validation and uncertainty characterization that were missing from the original manuscript and make explicit both the performance of the retrieval and the conditions under which its optical-property estimates become weakly constrained.

2.2 Systematic uncertainty due to stratospheric aerosol

Reviewer comment

A certain amount of aerosol is always to be expected in the stratosphere. When the SNR is low they can be hidden in the noise. They can introduce a negative systematic bias to the transmission method which depends on their scattering ratio.

How do the authors assess the systematic uncertainty introduced by stratospheric aerosol?

Response

We thank the reviewer for raising this important point. We agree that weak stratospheric aerosol contributions may remain below the cloud-detection threshold, particularly under low-SNR conditions, and may therefore introduce a systematic bias in a transmission-based COD estimate if the aerosol loading differs between the clear-air regions below and above the cloud.

The original scope of LACRIT was not to retrieve or characterize stratospheric aerosol. Nevertheless, the algorithm includes several safeguards that reduce sensitivity to non-molecular contributions. Cloud detection requires simultaneous consistency with both the profile-adaptive SNIP residual and the attenuated molecular reference, while the transmission retrieval is performed on the molecular-referenced ratio rather than directly on the elastic signal. Robust statistics are additionally used in the clear-air windows, and increased variability in these regions is propagated through σ_{COD} , thereby weakening the corresponding constraint on the lidar-ratio retrieval.

These procedures reduce the influence of isolated fluctuations and detectable particulate structures, but they do not guarantee the absence of weak, vertically extended aerosol that remains indistinguishable from the molecular background. We have therefore clarified this limitation in the revised manuscript. In particular, the synthetic validation presented here evaluates measurement-noise sensitivity down to a target SNR of 2, but it does not include an explicit stratospheric-aerosol component and should not be interpreted as a validation against aerosol-induced systematic bias.

To first order, the sensitivity of the molecular-referenced transmission retrieval to differential particulate contamination between the lower and upper reference windows can be expressed as

$$\Delta\text{COD}_{\text{aer}} \simeq -\frac{1}{2} \ln \left(\frac{R_{\text{aer,t}}}{R_{\text{aer,b}}} \right),$$

where $R_{\text{aer,b}}$ and $R_{\text{aer,t}}$ represent the effective particulate enhancement of the molecular-referenced signal in the lower and upper transmission windows. Thus, a larger undetected particulate contribution in the upper reference window produces a negative COD bias, whereas similar contributions in both windows largely cancel in the ratio.

We now explicitly state that aerosol contamination of the molecular reference and transmission windows remains a source of systematic uncertainty not fully quantified by the present synthetic experiment. A dedicated sensitivity study including prescribed stratospheric-aerosol layers would be required to characterize this contribution quantitatively.

3 Specific comments

3.1 Abstract terminology — Lines 5–8

Reviewer comment

This part of the abstract includes technical terms which might be used by some few experts on the field but it is not intuitive for general active remote sensing readers. Please consider improving the part by using less technical, more commonly recognizable terminology to make it clearer to a bigger audience.

Response

We thank the reviewer for this suggestion. We agree that the original abstract contained several method-specific terms that are unnecessarily technical for a general active-remote-sensing

readership. We have therefore simplified this part of the abstract while retaining the definitions of cloud optical depth (COD) and lidar ratio (LR), which are standard quantities in lidar studies.

In particular, expressions such as “robust log-median statistics”, “molecular reference windows”, and the detailed description of the particular-integration implementation have been replaced by a more physical description of the retrieval. The revised abstract now explains that COD is derived from the attenuation of the lidar signal across the cloud using clear-air molecular regions below and above the layer, while LR is obtained by combining this attenuation constraint with the enhancement of the elastic backscatter signal within the cloud.

The detailed mathematical formulation is retained in Sec. 3.

3.2 Lidar overlap and receiver configuration — Lines 90–91

Reviewer comment

Please mention here the distance of full overlap for each telescope.

Does the receiver after each telescope include the same kind of channels (e.g. both 532 p and s and 808)?

I would generally recommend to add a table with the information relevant to this study in the lidar section.

Response

We thank the reviewer for this suggestion. We have expanded the lidar description and the corresponding instrument-characteristics table to clarify both the receiver-channel configuration and the channel-dependent full-overlap distances.

The two receiving optical paths do not provide identical channel sets. The 532 nm receiver provides separate parallel and perpendicular polarization channels, whereas the second receiving path is dedicated to the 808 nm elastic channel. The full-overlap distances are approximately 1.275 km, 1.020 km, and 1.170 km for the 532 nm parallel, 532 nm perpendicular, and 808 nm channels, respectively.

Following the reviewer’s recommendation, these quantities have also been added to the instrument-characteristics table in the revised manuscript, which summarizes the instrumental parameters relevant to the present study.

3.3 Off-zenith pointing and specular reflection — Lines 107–108

Reviewer comment

Is the laser pointing at 5° off-zenith to avoid specular reflection?

Please include this information in the lidar description section. If not, how do you deal with specular reflections from cirrus clouds?

Response

We thank the reviewer for raising this point. The CE376 is not operated at a fixed 5° off-zenith angle specifically designed to suppress specular reflection. We have therefore clarified the actual pointing configuration in the lidar description and have avoided implying that specular reflection is systematically eliminated by the viewing geometry.

We acknowledge that near-zenith observations of cirrus may be affected by enhanced backscatter from horizontally oriented ice crystals. LACRIT does not currently include a dedicated correction for such specular-reflection events. Instead, the retrieval relies on the consistency of the elastic signal, depolarization information, cloud-layer structure, and ancillary observations used in the analysis. Cases showing anomalously enhanced backscatter associated with very low depolarization would therefore require particular caution.

We have added this clarification to the lidar description and now explicitly identify possible specular reflection from oriented ice crystals as a limitation of near-zenith cirrus observations.

3.4 Meaning of “absolute quantities” — Lines 131–132

Reviewer comment

It is not clear what “absolute quantities” refers to here. Can you provide some examples? (e.g. COD?)

Response

We thank the reviewer for pointing out this ambiguity. We agree that the expression “absolute quantities” was not appropriate and could be misinterpreted as referring to absolutely calibrated lidar products such as COD or LR.

What we intended to describe was the transition from the normalized, dimensionless signal quantities used throughout most of the processing chain to physically dimensioned optical quantities during the inversion. In particular, in the PI formulation the relative enhancement of the normalized elastic signal above the clear-air reference is multiplied by the molecular backscatter coefficient, providing a backscatter contribution with physical dimensions before conversion to an LR-dependent optical depth.

We have therefore replaced the expression “absolute quantities” with a more explicit description of this normalization-to-physical-quantity step.

3.5 Normalization and identification of molecular regions — Lines 128–133

Reviewer comment

Where does the normalisation take place? Can it be applied on both aerosol and molecular regions or on molecular regions only? Please specify.

Identifying molecular regions in a lidar signal is not a trivial process. The technique used to identify molecular regions must be explained here.

To what extent does the presence of aerosols (even in small, close to noise quantities) in the normalisation region affect your cloud retrievals?

Response

We thank the reviewer for these comments. We agree that the original description did not sufficiently distinguish the normalization step from the subsequent cloud-detection and inversion procedures.

In the revised manuscript, we now specify that normalization is performed only over a high-altitude molecular reference interval, selected within a molecular-dominated portion of the profile where particulate backscatter is expected to be negligible. The measured elastic range-corrected signal and the molecular reference profile are normalized over this same interval before being

used in the cloud-detection and transmittance calculations. The normalization is therefore not intentionally performed over aerosol-containing regions.

We have also clarified how the molecular reference region is selected. The analysis is restricted to an altitude interval above the boundary layer and below the region where the signal becomes too weak, and the normalization window is chosen within this domain where the measured profile is consistent with the expected molecular return and no significant particulate enhancement is identified. The subsequent SNIP-based background estimate, robust clear-air statistics, and molecular-consistency criterion provide additional checks against residual cloud or aerosol structures.

We nevertheless agree that the identification of a perfectly aerosol-free molecular region cannot be guaranteed from an elastic lidar profile alone. Weak aerosol loading close to the noise level may remain undetected and can bias the normalization and transmission estimate if its contribution differs between the reference regions used in the retrieval. This limitation is now stated explicitly in the revised manuscript.

As discussed in our response to the previous comment on stratospheric aerosol, the present synthetic validation quantifies sensitivity to measurement noise but does not include an explicit aerosol-background perturbation; therefore, it should not be interpreted as a complete assessment of this systematic uncertainty.

The complete revised formulation is provided in sec. 3.2, reproduced at the end of this response document.

3.6 Layering technique — Lines 134–139

Reviewer comment

This part is not clear. The authors have to either cite parts of the technique which have been already published or, if not published yet, they should describe the layering technique in more detail.

Please avoid using terms which are too ambiguous such as “residual hysteresis detection” and “adaptive thresholding” unless you describe them first.

In addition, consider linking this part to the corresponding subsection if a more detailed description is provided there.

Response

We thank the reviewer for this comment. We agree that the original methodology section introduced terms such as “residual hysteresis detection” and “adaptive thresholding” before these concepts had been formally defined, which made the description unnecessarily difficult to follow.

We have therefore reorganized Sec. 3 so that the introductory part now serves only as a roadmap of the LACRIT processing chain. Method-specific terminology has been removed from this overview, and each processing step is now explicitly linked to the subsection in which it is described in detail.

In particular, the revised cloud-detection subsection now defines the profile-adaptive SNIP background, the relative residual, the robust profile-specific noise estimate, the two-level hysteresis criterion, the molecular-consistency test, and the geometrical consolidation of detected layers. The numerical thresholds and persistence criteria are also summarized in a dedicated table.

This restructuring avoids introducing ambiguous terminology before its definition and provides a direct connection between the workflow shown in the methodology figure and the corresponding detailed methodological description.

The complete revised formulation is provided in sec. 3.2, reproduced at the end of this response document.

3.7 Consistency of molecular-profile terminology — Section 3

Reviewer comment

Many different terms such as molecular baseline, molecular backscatter, molecular reference are used to describe probably the same thing: the molecular attenuated backscatter profile.

Please avoid using multiple terms which introduce ambiguity, especially not clear ones such as molecular baseline, and stick to a single term and maybe an abbreviation for it for the whole manuscript.

For example in line 148, does β_{Ray} correspond to the molecular attenuated backscatter profile or to the molecular backscatter coefficient profile? These are two different things.

Response

We thank the reviewer for pointing out this ambiguity. We agree that the original manuscript used several terms such as “molecular baseline”, “molecular backscatter”, and “molecular reference” too loosely, although they refer to physically distinct quantities.

The terminology has therefore been revised throughout Sec. 3 to distinguish three quantities explicitly. First, $\beta_{\text{m}}(r)$ now denotes only the physical molecular backscatter coefficient computed from the thermodynamic profile. Second, $S_{\text{m}}(r)$, or $\tilde{S}_{\text{m}}(r)$ after normalization, denotes the expected attenuated molecular elastic return used as the molecular reference for detection and transmittance calculations. Third, the terms $B_{\text{SNIP}}(r)$ and $B_{\text{PI}}(r)$ are reserved for signal baselines estimated directly from clear-air portions of the measured profile and are no longer referred to as molecular baselines.

Accordingly, the quantity previously denoted $\beta_{\text{Ray}}(r)$ in the original manuscript corresponds to the molecular backscatter coefficient profile, not to the attenuated molecular backscatter signal. This notation has been replaced by $\beta_{\text{m}}(r)$ for consistency with the rest of the revised methodology.

The revised Sec. 3 now explicitly introduces these quantities at the stages where they are used and maintains the same terminology throughout the detection, transmission, and PI retrieval descriptions.

3.8 Sensitivity to temperature and pressure profiles — Lines 143–147

Reviewer comment

The transmission method is very sensitive to the molecular attenuated backscatter values before and after the cloud.

How much do you expect model errors in the temperature and pressure profiles to affect your retrievals?

The US standard atmosphere model can be quite inaccurate. Have you compared with radiosonde data to assess the introduced systematic errors?

Please also mention shortly which products you use from the microwave radiometer and how you apply them to calculate β_{Ray} .

Response

We thank the reviewer for this important comment. We agree that the transmittance retrieval depends on the accuracy of the molecular reference profile on both sides of the detected cloud and that uncertainties in the thermodynamic profiles may therefore propagate into the retrieved cloud optical depth.

Regarding the microwave-radiometer product and its use in the molecular backscatter calculation, in the present implementation the radiometer provides the temperature profile used to construct the molecular atmosphere over its useful altitude range, approximately up to 10 km. Above this altitude, the thermodynamic profile is completed using the Air Force Geophysics Laboratory (AFGL) standard atmosphere. The molecular backscatter coefficient is then calculated from the atmospheric number density derived from the pressure and temperature profiles and the molecular scattering properties at the lidar wavelength.

Regarding the sensitivity of the retrieval to uncertainties in the thermodynamic profile, following the reviewer’s comment, we evaluated the sensitivity of the molecular backscatter profile to the atmospheric data source. For a representative clear-sky case, the molecular backscatter coefficient derived using the AFGL atmospheric profile was compared with that obtained using the microwave-radiometer temperature information. Relative differences in β_{m} remained below approximately 4% over the considered altitude range.

It is important, however, that the transmission method is primarily sensitive to the differential error in the molecular reference between the clear-air regions below and above the cloud, rather than to a common multiplicative offset in β_{m} . If the relative molecular-reference errors in the lower and upper windows are denoted by ϵ_b and ϵ_t , respectively, their first-order contribution to the retrieved optical depth is

$$\Delta\text{COD}_{\text{trans}} \simeq \frac{1}{2} (\epsilon_t - \epsilon_b).$$

Consequently, thermodynamic errors that vary slowly with altitude tend to partially cancel in the differential transmission estimate, whereas vertical variations in the thermodynamic error across the cloud altitude range may introduce a systematic COD bias.

Regarding the reviewer’s suggestion to assess this uncertainty using radiosonde observations, a dedicated radiosonde intercomparison was not performed for the present study. The sensitivity assessment introduced in this revision instead compares the molecular backscatter obtained from the microwave-radiometer information with that derived from the AFGL atmospheric profile. We have made this limitation explicit in the revised manuscript and avoid implying that a radiosonde-based validation was performed.

The revised manuscript has accordingly been updated to clarify the microwave-radiometer product used in the molecular-profile construction, the use of the modeled atmosphere above the radiometer range, and the sensitivity of the molecular backscatter and transmission-based COD retrieval to the adopted thermodynamic profile.

3.9 Lidar-provided quantities and wavelength — Lines 148–149

Reviewer comment

Please mention that RCS and VLDR are provided by the lidar and also mention the wavelength (532 nm).

Response

We thank the reviewer for this suggestion. The paragraph originally referred to in Lines 148–149 has been substantially reorganized as part of the revision of Sec. 3. In the revised manuscript, the lidar-provided quantities are now introduced explicitly in the data-preprocessing subsection.

The text now states that the elastic range-corrected signal (RCS) and the volume linear depolarization ratio (VLDR) are provided by the CE376 lidar, and that the present analysis is performed using the polarized 532 nm channel. This information is therefore now incorporated directly into the revised methodological description rather than being retained in the original paragraph.

The complete revised formulation is provided in sec. 3.1, reproduced at the end of this response document.

3.10 Definition of the cloud core — Lines 164–168

Reviewer comment

How do you identify the cloud core region? How big is the window?

Response

We thank the reviewer for this comment. We have clarified that two distinct vertical regions are involved in the processing: the cloud core used for automatic layer detection and the sampling window used to derive a representative VLDR value.

The cloud-detection algorithm first identifies statistically significant core bins satisfying both the SNIP-residual and molecular-reference criteria. Once a core is established, the detected region is extended toward weaker cloud edges using the lower hysteresis thresholds together with persistence and geometrical criteria, as described in Sec. 3.2 and illustrated in the Appendix. In the configuration adopted for the synthetic validation, candidate cores thinner than 45 m are rejected and final detected layers thinner than 150 m are discarded. These values are configurable quality-control parameters used to stabilize the automatic detection and do not represent physical limits on cirrus-cloud thickness.

The VLDR sampling region is defined separately from these detection criteria. After the final layer geometry has been determined, the representative VLDR is evaluated within a configurable vertical interval centred on the retrieved signal peak. This interval is independent of both the initial detection core and the final cloud boundaries, which limits the contribution of molecular background and weak edge bins to the representative depolarization value. The corresponding description has been revised in Sec. 3.2 to make this distinction explicit.

The complete revised formulation is provided in sec. 3.2 and Appendix A3, reproduced at the end of this response document.

3.11 Automatic identification of aerosol-free regions — Lines 169–171

Reviewer comment

Automatic identification of aerosol-free regions is not a trivial process. Please either cite or describe the applied technique.

Response

We thank the reviewer for this comment. We agree that the automatic identification of truly aerosol-free regions cannot be assumed without further description.

The revised Sec. 3 now explains the procedure used to identify the molecular-dominated reference regions employed in the normalization and transmittance calculations. The analysis is first restricted to an altitude range above the boundary layer and below the region where the lidar signal becomes too weak. Within this domain, the normalization interval is selected where no significant particulate enhancement is identified and where the measured elastic signal remains consistent with the expected attenuated molecular return.

The revised cloud-detection subsection further describes how the measured profile is evaluated relative to both a profile-adaptive SNIP background and the normalized attenuated molecular reference, with robust clear-air statistics used to reduce the influence of isolated residual structures. These criteria are described explicitly in Sec.3.2, and the corresponding detection logic and configuration parameters are illustrated and documented in the Appendix A3.

We have also clarified that these procedures identify molecular-dominated rather than demonstrably aerosol-free regions. Weak particulate contributions close to the noise level may remain undetected, and this residual contamination is now acknowledged as a potential source of systematic uncertainty.

The complete revised formulation is provided in sec. 3.2 and Appendix A3, reproduced at the end of this response document.

3.12 Role of SNIP — Lines 177–178

Reviewer comment

The molecular attenuated backscatter profile without attenuation only from molecules can be already calculated from the temperature and pressure profile and the molecular backscatter and extinction cross-sections.

Please explain how SNIP is applied and what does it additionally offer.

Response

We thank the reviewer for this clarification. We agree that the attenuated molecular reference can be calculated directly from the temperature and pressure profiles together with the molecular backscatter and extinction properties, and that SNIP should not be interpreted as an alternative method for constructing this molecular profile.

In the revised methodology, these two quantities are now explicitly separated. The attenuated molecular reference $\tilde{S}_m(r)$ is calculated from the thermodynamic molecular profile and is used as a physical consistency reference. By contrast, SNIP is applied directly to the normalized measured elastic profile $\tilde{S}(r)$ in order to estimate a smooth, profile-adaptive background

$B_{\text{SNIP}}(r)$ after iteratively suppressing positive local enhancements associated with cloud-like structures.

The additional role of SNIP is therefore not to reproduce the molecular atmosphere, but to provide a local empirical background against which statistically significant signal enhancements can be detected. This is useful because the measured profile may contain slowly varying instrumental effects, residual background structure, or weak non-molecular contributions that are not represented by the ideal molecular reference alone.

LACRIT consequently uses the two quantities in a complementary way: $B_{\text{SNIP}}(r)$ determines whether a localized enhancement is significant relative to the measured-profile background, whereas $\tilde{S}_m(r)$ determines whether the same enhancement also exceeds the return expected from a molecular atmosphere. We have revised Sec. 3.2 and Appendix A3 to make this distinction explicit.

3.13 Definition of residual statistics, detection thresholds, and transmission estimator — Lines 188–205

Reviewer comment

Median of what? Do you refer to the residuals? In which region are the residuals calculated?

Information regarding how the low and high thresholds are defined is missing here. Are the thresholds applied to the residuals? This is also not clear. This part has to be written better.

The term robust log-median statistics is not self explanatory. Please describe why the applied method is better than, e.g., just calculating the two-way transmittance by averaging S inside the two windows W_t and W_b and then dividing them.

Response

We thank the reviewer for these related comments. We agree that the original description did not clearly define the residual quantity, the region over which its statistics were evaluated, the construction of the high and low detection thresholds, or the meaning of the expression “robust log-median statistics”. This part of Sec. 3 has therefore been completely rewritten.

In the revised cloud-detection subsection, the detection variable is explicitly defined as the relative SNIP residual

$$R_{\text{SNIP}}(r) = \frac{\tilde{S}(r) - B_{\text{SNIP}}(r)}{|B_{\text{SNIP}}(r)| + \varepsilon},$$

which is evaluated over the molecular-dominated altitude range considered by the detector. The effective background variability is then estimated from clear-air portions of this residual using

$$\sigma_R = 1.4826 \text{ MAD}(R_{\text{SNIP,clear}}),$$

after iterative clipping of residual cloud-related excursions. Thus, the “median” referred to in the original text is no longer used ambiguously; the revised manuscript now specifies explicitly whether a median is applied to the residual distribution or to the molecular-referenced transmission quantity.

The two hysteresis thresholds are applied directly to R_{SNIP} . A restrictive core threshold,

$$R_{\text{core}} = \max(0.020, 5\sigma_R),$$

is used to establish statistically significant cloud cores, whereas the lower

$$R_{\text{grow}} = \max(0.003, 1.5\sigma_R)$$

threshold is used to extend those cores toward weaker cloud edges. The corresponding molecular-consistency margins are applied simultaneously using the normalized attenuated molecular reference.

The transmission calculation has also been rewritten to avoid the ambiguous term “robust log-median statistics”. After cloud detection, LACRIT forms the molecular-referenced ratio

$$q(r) = \frac{\tilde{S}(r)}{\tilde{S}_m(r)}$$

and calculates the cloud optical depth from the difference between the median values of $\ln q(r)$ in the clear-air windows below and above the cloud:

$$\text{COD}_{\text{trans}} = -\frac{1}{2} [\text{median}_{W_t}(\ln q) - \text{median}_{W_b}(\ln q)].$$

This formulation was selected for robustness rather than because averaging the signals is intrinsically incorrect. Taking the logarithm converts the multiplicative transmission ratio into an additive difference, while using the median reduces sensitivity to isolated spikes, residual cloud-edge contamination, and non-Gaussian fluctuations within the reference windows. A simple ratio of mean signals would be more sensitive to a small number of anomalously large bins. We have revised the text to state this rationale explicitly and to avoid implying that the log-median estimator is universally superior in all noise regimes.

These definitions are now introduced sequentially in Sec.3.2, so that the residual statistics, threshold construction, and transmission estimator are described within a single consistent methodological framework.

The complete revised formulation is provided in sec. 3.2, reproduced at the end of this response document.

3.14 Offset addition and lifted signal — Lines 213–217

Reviewer comment

It is not clear neither what the physical meaning of the offset addition is nor how transmittance is recomputed using the so-called lifted signal.

The transmission method is rather straightforward, it relies on finding a pure molecular region before and after the cloud. If there is aerosol contamination in one of those regions then the result will be biased.

It is not possible to “tune” the result without having a priori information about the amount of aerosol contamination. At best, one can assume a conservative amount of aerosols based on the RCS and its noise and use it for the error simulation.

How do you control how much to “lift” the signal and when does the lifting stop?

Response

We thank the reviewer for this important comment. We agree that the original description of the “lifted signal” was insufficiently clear and could be misinterpreted as an empirical adjustment intended to compensate for aerosol contamination.

The offset-correction branch shown in the LACRIT workflow is not part of the nominal retrieval used for the observational cases presented in this study. These cases were processed using the standard transmittance configuration, without applying an additive signal correction. The `OFFSET_TRANSM` branch is retained in the processing framework only as an optional fallback for particularly difficult profiles in which the standard transmittance calculation fails its quality-control criteria because of poor signal conditions or an apparent additive baseline displacement.

When this optional branch is activated, the correction is not iteratively tuned. The standard transmittance retrieval is attempted first. If it fails, a sufficiently extended clear-air region outside the detected cloud is identified and a single additive correction is estimated from the median difference between the normalized attenuated molecular reference and the normalized measured signal,

$$c = \text{median}_{W_{\text{clear}}} \left[\tilde{S}_m(r) - \tilde{S}(r) \right].$$

The corrected signal is then

$$\tilde{S}_{\text{corr}}(r) = \tilde{S}(r) + c,$$

and the transmittance calculation is repeated once. The correction is accepted only when its magnitude remains below a predefined limit, and the resulting transmittance must still satisfy the standard positivity, SNR, and physical consistency criteria. There is therefore no iterative “lifting” toward a prescribed COD value.

We fully agree with the reviewer that this procedure cannot determine or correct an unknown aerosol contribution in the transmission reference windows. Residual aerosol contamination is treated separately as a potential source of systematic uncertainty. To avoid ambiguity, the revised manuscript now presents the offset correction only as an optional fallback branch of the processing framework and explicitly states that it was not used for the study cases analyzed here.

3.15 Temporal regularization — Lines 237–246 and Eq. 9

Reviewer comment

How long are the temporal intervals when you apply temporal regularization? Is it minutes, seconds, or hours?

Response

We thank the reviewer for this comment. We agree that the temporal scale of the regularization was not specified clearly in the original manuscript.

The temporal regularization is applied between successive lidar retrievals, so its effective time interval is set by the temporal spacing of the processed profiles. For the observational analyses presented in this study, the lidar profiles are typically averaged over 1–5 min, and the temporal constraint, when used, therefore acts over consecutive retrievals at that cadence.

For the synthetic temporal-sensitivity experiment, a fixed cadence of 120 s between successive profiles was prescribed in order to evaluate the behaviour of the temporal constraint under controlled conditions. The synthetic sequence contains 12 profiles.

We have added these temporal scales explicitly to the revised methodology and validation description.

3.16 Uncertainty assessment and influence of the lidar-ratio prior — Lines 247–269

Reviewer comment

This part is not sufficient to explain a rather complicated uncertainty schema. It is not clear to me how the uncertainty is currently assessed.

It is rather important to quantify well the uncertainty of the technique because many optimization methods are applied and classical propagation of uncertainty (Monte Carlo or analytical) is no longer possible.

How can you exclude that the low variability of the lidar ratio is not due to the a-priori regularization?

How much is the applied a-priori lidar ratio?

Fig. 4 does not show any information about the lidar ratio, were you referring to Table 1 in the text?

Response

We thank the reviewer for these related comments. We agree that the original manuscript did not sufficiently distinguish the observational uncertainty of the retrieval from the stabilizing effect of the optional regularization terms. This part of the methodology and validation has therefore been substantially expanded.

In the revised formulation, the uncertainty of the transmission-based optical depth is estimated directly from the variability of the molecular-referenced signal in the clear-air windows below and above the cloud,

$$\sigma_{\text{COD}} = \frac{1}{2} \sqrt{\sigma_{\ln q, b}^2 + \sigma_{\ln q, t}^2},$$

and this uncertainty enters explicitly into the lidar-ratio objective function,

$$J_{\text{data}}(L) = \left[\frac{\text{COD}_{\text{PI}}(L) - \text{COD}_{\text{trans}}}{\sigma_{\text{COD}}} \right]^2.$$

Thus, profiles for which the transmission constraint is poorly determined automatically provide a weaker observational constraint on LR.

Because the complete retrieval includes non-linear optimization and optional regularization, we agree that a simple analytical propagation is not sufficient to characterize the total retrieval behavior. We therefore introduced an independent synthetic Monte Carlo validation in which prescribed COD and LR values are used to generate elastic-lidar profiles over a controlled range of SNR conditions. Fifty independent noise realizations are processed for each finite-SNR COD–LR combination, allowing the resulting distributions of detection, COD, and LR errors to be quantified directly.

To isolate the influence of the prior, the synthetic experiment explicitly compares the unregularized `DATA_ONLY` retrieval with a `PRIOR_ONLY` configuration using exactly the same synthetic profiles and paired noise realizations. The prior used in this sensitivity experiment is deliberately broad and is imposed independently of the prescribed truth, with $L_0 = 45$ sr and a width of 25 sr. It is therefore not intended as a climatological estimate of cirrus LR, but as a controlled regularization experiment.

The comparison shows that the prior reduces LR dispersion and RMSE as the observational constraint weakens, but also introduces the expected shrinkage toward L_0 . We therefore do not interpret the reduced LR variability in the regularized retrieval as independent evidence of atmospheric stability. The unregularized and regularized results are now reported separately, so

that the effect of the prior can be identified explicitly rather than being conflated with retrieval precision.

To isolate the influence of the different regularization terms, the synthetic validation considers the retrieval configurations summarized in table 4 which has been included in the manuscript (Table 4).

Table 1: Retrieval configurations used in the synthetic validation.

Configuration	Objective function	Additional constraint
DATA_ONLY	J_{data}	None
PRIOR_ONLY	$J_{\text{data}} + J_{\text{prior}}$	Static LR prior
TEMPORAL_ONLY	$J_{\text{data}} + J_{\text{time}}$	Temporal continuity
PRIOR+TEMPORAL	$J_{\text{data}} + J_{\text{prior}} + J_{\text{time}}$	Static prior + temporal continuity

We have also corrected the erroneous reference to Fig. 4 in the original text. The lidar-ratio information referred to there was associated in the corresponding retrieval-results figure and table.

The complete revised formulation is provided in sec. 3.3 and Appendix B–D, reproduced at the end of this response document.

4 Technical comments

4.1 Typographical correction — Lines 107–108

Reviewer comment

Please correct the typo “were serve”.

Response

We thank the reviewer for pointing this out. The typographical error “were serve” has been corrected in the revised manuscript.

4.2 Figure 4

Reviewer comment

I don’t see the red dashed line (baseline), is it overlapped?
It is not clear why the cyan line (polynomial fit) is there and what does it show.

Response

We thank the reviewer for this comment. The red dashed SNIP baseline is indeed present in Fig. 4, but over most of the profile it nearly overlaps with the polynomial clear-air fit and the normalized Rayleigh reference. Because the horizontal scale is dominated by the much larger cloud backscatter peaks, the small differences among these reference curves are visually compressed.

The cyan curve represents the low-order polynomial fit constructed from clear-air portions of the measured profile. In LACRIT, this polynomial fit is used preferentially as the clear-air reference for the modified PI retrieval, whereas the SNIP baseline is primarily used for cloud detection and can also serve as an alternative PI reference when a stable polynomial fit is not available. The normalized Rayleigh reference is used for the molecular consistency test and for the transmission-based COD retrieval.

We have clarified these respective roles in the revised figure caption.

4.3 Meaning of “post-processed” — Line 258

Reviewer comment

What does “post-processed” here mean? Is additional processing other than LACRIT applied?

Response

We thank the reviewer for pointing this out. The term “post-processed” was unnecessary and potentially misleading, since no additional processing beyond the LACRIT workflow was intended. The sentence referred only to the use and analysis of the output products generated by the algorithm. We have therefore removed the term from the revised manuscript.

4.4 Figure 6

Reviewer comment

Which of the two cases does this figure correspond to?

Response

We thank the reviewer for this comment. Figure 6 corresponds to the first case study, observed on 25 December 2022. This has now been specified in the figure caption.

4.5 Figures 5 and 7

Reviewer comment

Please mark Layer 1 and Layer 2 on the plots.

Response

We thank the reviewer for this suggestion. Layer 1 and Layer 2 are now explicitly identified in the corresponding lower panels of the revised figure. We have kept the time–height RCS panel unchanged in order to avoid overloading the figure with additional annotations.

Appendix to the Response:

Substantially revised sections

3.1 Data preprocessing

The retrieval is based on elastic lidar measurements described by the standard single-scattering lidar equation,

$$P(R) = \frac{K O(R)}{R^2} \beta(R) \exp \left[-2 \int_0^R \alpha(r) dr \right], \quad (1)$$

where $P(R)$ is the received power at range R , K is the system constant, $O(R)$ is the overlap function, and $\beta(R)$ and $\alpha(R)$ denote the total volume backscatter and extinction coefficients, respectively, including molecular and particulate contributions. The elastic observations are expressed as range-corrected signals, $S(R) = P(R)R^2$, which removes the geometrical R^{-2} dependence and facilitates comparison of atmospheric structures at different altitudes.

In parallel with the measured elastic signal, a molecular atmospheric profile is constructed on the lidar altitude grid. The molecular backscatter coefficient, $\beta_m(R)$, is derived from the thermodynamic state of the atmosphere. When available, microwave-radiometer observations are used within their useful altitude range, extending to approximately 10 km. Above this altitude, or when radiometer observations are unavailable, the thermodynamic profile is completed using the best available ancillary atmospheric information, such as nearby atmospheric soundings or modelled/standard atmospheric profiles. The resulting temperature, pressure, and humidity fields are used to derive atmospheric number density and the corresponding molecular backscatter and extinction coefficients. Throughout the manuscript, $\beta_m(R)$ denotes the physical molecular backscatter coefficient.

Before automatic layer detection, the lidar profiles are temporally averaged (typically 1–5 min), vertically smoothed, and restricted to an altitude domain above the atmospheric boundary layer where the molecular contribution is expected to dominate. This restriction reduces overlap-related effects and limits the influence of the stronger aerosol loading generally encountered at lower altitudes.

The molecular backscatter and extinction profiles are then combined to construct the attenuated molecular elastic return $S_m(R)$ expected for the same atmospheric state. This quantity is distinct from $\beta_m(R)$: the latter is a physical backscatter coefficient, whereas $S_m(R)$ represents the altitude-dependent elastic molecular return after molecular attenuation.

To compare the measured and molecular signals on the same relative scale, a high-altitude reference interval $[r_1, r_2]$ is selected within a molecular-dominated region where no significant particulate enhancement is identified and where particulate backscatter is expected to remain small relative to the molecular return. The measured profile is normalized according to

$$\tilde{S}(R) = \frac{S(R)}{\langle S(r) \rangle_{r \in [r_1, r_2]}}, \quad (2)$$

and the attenuated molecular signal is normalized over the same interval,

$$\tilde{S}_m(R) = \frac{S_m(R)}{\langle S_m(r) \rangle_{r \in [r_1, r_2]}}. \quad (3)$$

This common normalization removes the unknown multiplicative system response from the principal relative quantities used by LACRIT and therefore avoids the need for an absolute radiometric calibration of the elastic signal. The normalized measured and molecular profiles, $\tilde{S}(R)$ and $\tilde{S}_m(R)$, constitute the two principal signal references used in the subsequent detection and transmittance calculations. Physical dimensions are introduced later in the radiative-property

retrieval through the molecular backscatter coefficient $\beta_m(R)$.

3.2 Cloud-layer detection and molecular-referenced transmittance

Cloud detection begins by identifying localized enhancements of the normalized elastic signal relative to its slowly varying measured background. Rather than assuming a fixed analytical form for this background, LACRIT estimates it directly from each profile $\tilde{S}(r)$ using the Statistics-sensitive Non-linear Iterative Peak-clipping (SNIP) algorithm.

SNIP iteratively suppresses localized positive excursions by comparing each signal bin with its surrounding neighbourhood over progressively varying window sizes, thereby retaining the slowly varying component of the measured profile while removing cloud-like peaks from the baseline estimate.

The resulting data-derived baseline, $B_{\text{SNIP}}(r)$, therefore represents the local measured background against which potential cloud enhancements are evaluated.

Cloud-related departures from this baseline are quantified through the relative residual

$$R_{\text{SNIP}}(r) = \frac{\tilde{S}(r) - B_{\text{SNIP}}(r)}{|B_{\text{SNIP}}(r)| + \varepsilon}, \quad (4)$$

where ε is used only to prevent numerical instability when the baseline approaches very small values. The residual therefore expresses the strength of a local signal enhancement relative to the measured atmospheric background rather than in absolute signal units.

The statistical significance of each residual enhancement is evaluated relative to the variability of the clear-air residual itself. For each profile, an effective residual-noise level is estimated as $\sigma_R = 1.4826 \text{ MAD}(R_{\text{SNIP,clear}})$, with iterative $3\sigma_R$ clipping applied to reduce the influence of residual cloud-related excursions. The resulting σ_R therefore provides a profile-specific measure of background variability from which the adaptive detection thresholds are constructed.

These thresholds are applied within a two-level hysteresis scheme. A restrictive threshold is first used to identify statistically significant cloud-core bins, while a lower threshold subsequently allows each established core to grow toward weaker cloud edges. In this way, the higher threshold controls the initial detection significance, whereas the lower threshold preserves vertically extended or optically weak portions of the same structure without allowing isolated noise fluctuations to develop into independent layers.

Statistical significance relative to the SNIP background is, however, not sufficient by itself to identify a cloud. A positive residual only indicates that the measured signal exceeds its local measured background; similar enhancements may also arise from weak particulate contributions or imperfections in the empirical baseline. LACRIT therefore complements the SNIP criterion with a second, physically based comparison against the normalized attenuated molecular reference. The corresponding molecular excess is defined as

$$R_{\text{mol}}(r) = \frac{\tilde{S}(r)}{\tilde{S}_m(r) + \varepsilon} - 1. \quad (5)$$

The two residuals thus provide complementary information: R_{SNIP} quantifies the significance of a feature relative to the observed local background, whereas R_{mol} evaluates whether the same feature also exceeds the expected attenuated molecular return.

A cloud core is retained only when the required SNIP and molecular criteria are simultaneously satisfied. The detection is then expanded toward weaker cloud edges using the lower hysteresis thresholds, after which persistence and geometrical criteria are applied to suppress isolated structures, constrain boundary growth, enforce a minimum vertical extent, and merge nearby detections. This procedure yields the final layer geometry, including cloud base r_b , cloud top r_t , and the altitude of maximum residual enhancement.

The retrieved peak altitude is subsequently used to define a representative δ_v sampling region for cloud characterization. δ_v is evaluated within a configurable interval centred on the signal

peak ($\text{peak} \pm \Delta r$), independently of both the initial hysteresis core and the final layer boundaries. This choice reduces the contribution of weak edge bins and surrounding molecular background to the representative depolarization value.

The numerical settings governing layer detection and boundary determination are summarized in Table 2. The values reported here correspond to the configuration adopted for the synthetic validation and represent algorithmic quality-control and persistence settings rather than physical definitions of cirrus-cloud thickness. They may therefore be adapted to other instrumental configurations or applications.

Table 2: Cloud-detection parameters adopted for the synthetic validation of LACRIT.

Parameter	Value
SNIP half-window candidates	200, 300, 400, 500, 650, 800 m
Core threshold	$\max(0.020, 5\sigma_R)$
Growth threshold	$\max(0.003, 1.5\sigma_R)$
Minimum peak significance	$5\sigma_R$
Molecular core margin	0.020
Molecular growth margin	0.003
Core minimum thickness	45 m
Final minimum layer thickness	150 m
Persistence for boundary growth	4 bins
Merge distance	90 m
Robust clipping	$3\sigma_R$, 3 iterations

Once the layer boundaries have been established, its optical attenuation is estimated from clear-air windows W_b and W_t located below and above the cloud, respectively. A direct comparison of the normalized elastic signal in these two windows would retain the natural altitude dependence of the molecular return. To remove this dependence, LACRIT first forms the molecular-referenced ratio

$$q(r) = \frac{\tilde{S}(r)}{\tilde{S}_m(r)}, \quad (6)$$

so that variations associated with the expected molecular profile are largely removed before the cloud transmission is evaluated.

The transmission-based cloud optical depth is then obtained from the difference between the logarithmic signal levels in the upper and lower reference windows,

$$\text{COD}_{\text{trans}} = -\frac{1}{2} [\text{median}_{W_t} (\ln q(r)) - \text{median}_{W_b} (\ln q(r))]. \quad (7)$$

The logarithmic formulation converts the multiplicative two-way transmission into an additive difference between the two windows, while the median reduces sensitivity to isolated signal excursions, residual cloud-edge contamination, and non-Gaussian fluctuations. The factor $1/2$ accounts for the two-way propagation of the lidar signal through the layer.

The variability of $\ln q$ within the same windows is used to quantify the uncertainty of the transmission constraint,

$$\sigma_{\text{COD}} = \frac{1}{2} \sqrt{\sigma_{\ln q, b}^2 + \sigma_{\ln q, t}^2}, \quad (8)$$

where $\sigma_{\ln q, b}$ and $\sigma_{\ln q, t}$ describe the variability in the lower and upper windows, respectively. A larger variability therefore produces a larger σ_{COD} and, as described in sec. 3.3, a weaker transmission constraint on the subsequent lidar-ratio retrieval.

This molecular referencing corrects for the expected vertical variation of the clear-air molecular return, but it does not guarantee that the two reference windows are completely free of particulate contributions. In particular, different aerosol contributions below and above the cloud may introduce a systematic bias in $\text{COD}_{\text{trans}}$. Such contamination cannot be uniquely inferred from the elastic signal alone and is therefore treated as a limitation of the transmission method rather than as a quantity that can be empirically corrected.

3.3 Lidar-ratio retrieval and regularization

The transmittance method provides a direct constraint on the integrated attenuation produced by the cloud, but it does not by itself separate particulate extinction from backscatter and therefore cannot determine the lidar ratio. To introduce sensitivity to LR, LACRIT uses a modified implementation of the Particular Integration (PI) method of Cadet et al. 2005. The two approaches exploit different characteristics of the elastic signal: the transmittance method quantifies the loss of signal across the cloud, whereas the PI formulation relates the enhancement of the in-cloud return above a clear-air reference to the particulate backscatter.

For the PI calculation, the clear-air signal that would be expected in the absence of the detected cloud is represented by a data-derived clear-air reference $B_{\text{PI}}(r)$ reconstructed from regions outside the layer. In the present implementation, this reference is obtained preferentially from a second-degree polynomial fitted to the available clear-air bins, while the SNIP estimate provides an alternative when a stable polynomial reference cannot be constructed. A lower bound relative to the normalized attenuated molecular reference signal $\tilde{S}_m(r)$ is additionally imposed to prevent the reference signal from becoming artificially small, which would otherwise amplify the inferred particulate contribution.

For a trial particulate lidar ratio L , the PI estimate of cloud optical depth is written as

$$\text{COD}_{\text{PI}}(L) = L \int_{r_b}^{r_t} \left[\frac{\tilde{S}(r)}{B_{\text{PI}}(r)} - 1 \right] \beta_m(r) dr, \quad (9)$$

where $\beta_m(r)$ is the physical molecular backscatter coefficient evaluated on the same altitude level. The quantity in brackets represents the relative enhancement of the measured elastic return with respect to the estimated clear-air background, while multiplication by $\beta_m(r)$ converts this relative enhancement into a backscatter contribution with the appropriate physical dimensions. The dependence on L then transforms the inferred particulate backscatter into extinction and, after vertical integration, into an LR-dependent optical depth.

The lidar ratio can therefore be constrained by requiring consistency between two estimates of the same integrated cloud extinction obtained from different parts of the lidar signal. The transmittance estimate, $\text{COD}_{\text{trans}}$, is determined from the attenuation observed across the layer, whereas $\text{COD}_{\text{PI}}(L)$ depends on the in-cloud signal enhancement for a prescribed trial LR. In the absence of additional regularization, the retrieved LR corresponds to the value of L that minimizes

$$J_{\text{data}}(L) = \left[\frac{\text{COD}_{\text{PI}}(L) - \text{COD}_{\text{trans}}}{\sigma_{\text{COD}}} \right]^2, \quad (10)$$

The denominator carries the uncertainty of the transmission estimate into the inversion. Consequently, agreement between the two COD estimates provides a strong constraint when the clear-air transmission is well determined, whereas the cost function becomes progressively less discriminating when σ_{COD} increases. This distinction is important because the retrieval is not based simply on forcing two COD values to agree: the strength of that agreement is explicitly weighted according to the information content of the transmission measurement. The numerical settings defining the transmission windows, PI reference, and LR search space are summarized in Table 3.

Table 3: Radiative-property retrieval settings adopted for the synthetic validation of LACRIT.

Parameter	Value
Transmission-window thickness	200 m
Boundary-to-window separation	60 m
Transmission smoothing	9 bins
Minimum transmission SNR	1
Minimum σ_{COD}	10^{-4}
PI baseline	Second-degree polynomial
Alternative PI baseline	SNIP
PI baseline lower bound	$0.6 \tilde{S}_m$
LR search interval	10–100 sr
LR grid spacing	1 sr

Although the LR retrieval adopted in this study is based on the modified PI–transmittance constraint described above, LACRIT also computes forward and backward Fernald–Klett solutions as complementary consistency diagnostics. These inversions provide independent LR-dependent reconstructions of the elastic signal and are used to assess the behaviour of the retrieved solution, but they do not determine the LR values reported in the present analysis.

When the transmission constraint becomes weak, however, a broad range of LR values may produce similarly acceptable values of J_{data} . LACRIT therefore allows independent information to be introduced as soft regularization rather than as a prescribed LR solution. A static prior adds the term

$$J_{\text{prior}}(L) = \lambda_{\text{prior}} \left[\frac{\ln L - \ln L_0}{\sigma_{\ln L}} \right]^2, \quad (11)$$

where L_0 and $\sigma_{\ln L}$ describe the centre and width of the prior distribution. The prior is formulated in logarithmic LR space, such that $\ln L - \ln L_0 = \ln(L/L_0)$. The regularization therefore penalizes relative rather than absolute departures from the prior, treating equal multiplicative deviations above and below L_0 symmetrically. Importantly, the information entering this constraint does not need to originate from an LR climatology alone. More generally, the prior parameters may be obtained by independent observations or classifications associated with the scattering state of the cloud.

The static prior is not the only source of additional information that can be used when the instantaneous lidar constraint becomes weak. If consecutive retrievals are available, LACRIT can instead, or additionally, exploit the expected temporal continuity of the same cloud layer. This is introduced through

$$J_{\text{time}}(L_t) = \lambda_{\text{time}} \left[\frac{\ln L_t - \ln L_{t-1}}{\sigma_{\text{time}}} \right]^2, \quad (12)$$

where L_{t-1} is the LR retrieved at the preceding time step. The temporal term penalizes variations that are large compared with the assumed temporal variability of the layer and therefore acts as a continuity constraint rather than as an external LR estimate. Its usefulness depends on the retrieval cadence and on the rate at which the cloud properties evolve.

The temporal interval entering this constraint is set by the separation between successive processed lidar profiles. In the observational processing, profiles are typically averaged over 1–5 min; when temporal regularization is enabled, the constraint therefore operates between consecutive retrievals at that cadence. In the controlled synthetic experiment, a fixed temporal spacing of 120 s is used.

Taken together, the retrieval may therefore combine three sources of information: the radiative constraint provided by the current lidar profile, an optional external LR prior, and an optional temporal-continuity constraint derived from the preceding retrieval. The latter two terms act only as regularization: they do not constitute additional measurements of LR, but restrict the range of admissible solutions when the information contained in the elastic signal becomes insufficient. Their contribution must consequently be distinguished from the intrinsic performance of the unregularized retrieval.

This distinction is explicitly examined in the synthetic sensitivity analysis described in Sec. 4 and Appendix C, where the effects of the static and temporal regularization terms are evaluated separately from the data-only retrieval. The static-prior experiment uses a deliberately broad prior centred at $L_0 = 45$ sr with a width of 25 sr, independently of the prescribed synthetic LR, so that its influence can be quantified without providing prior knowledge of the true solution.

4 Results and discussion

4.1 Synthetic validation of LACRIT

To isolate the influence of the different regularization terms, the synthetic validation considers the retrieval configurations summarized in Table 4. In all cases, the observational term J_{data} is retained; the configuration names indicate which additional regularization terms are activated.

Table 4: Retrieval configurations used in the synthetic validation.

Configuration	Objective function	Additional constraint
DATA_ONLY	J_{data}	None
PRIOR_ONLY	$J_{\text{data}} + J_{\text{prior}}$	Static LR prior
TEMPORAL_ONLY	$J_{\text{data}} + J_{\text{time}}$	Temporal continuity
PRIOR+TEMPORAL	$J_{\text{data}} + J_{\text{prior}} + J_{\text{time}}$	Static prior + temporal continuity

The synthetic validation first examines the intrinsic behavior of the retrieval under deterministic conditions, before progressively degrading the signal to assess its sensitivity to measurement noise. In the absence of noise, Fig. 1a shows that the retrieved LR generally follows the prescribed values over the complete COD–LR matrix, with the best agreement obtained for optically thin cases. A systematic positive departure develops, however, as COD increases. This trend is summarized in Fig. 1b, where the mean LR bias evolves from -0.5 sr at $\text{COD} = 0.02$ to $+10.5$ sr at $\text{COD} = 0.60$. Across the complete noise-free matrix, the DATA_ONLY LR RMSE is 5.54 sr.

Because the synthetic forward model contains neither random measurement noise nor multiple-scattering effects, the residual increase in LR bias at high COD cannot be attributed to either of these processes. Instead, it identifies an intrinsic limitation of the present PI-based retrieval under the assumptions of the synthetic framework. The deterministic experiment therefore establishes the baseline algorithmic accuracy against which the subsequent noise-dependent behavior can be interpreted.

When measurement noise is introduced, the validation reveals a clear distinction between identifying the cloud and constraining its optical properties. Cloud detection remains highly robust over the tested SNR range. As shown in Fig. 2a, truth-matched detections remain at 100 % down to SNR 5 and decrease only slightly, to 99.4 %, at SNR 2. Similarly, the more restrictive requirement that both cloud boundaries lie within 150 m of their prescribed positions remains above 98 % for all tested noise levels.

The effect of decreasing SNR is more evident in the geometrical and radiative retrieval errors. Cloud-base and cloud-top RMSE increase markedly only at the lowest SNR, reaching approximately 130 and 122 m, respectively (Fig. 2b). In contrast, the DATA_ONLY LR-error distribution broadens progressively as the measurement constraint weakens (Fig. 2c). The syn-




Figure 1: Noise-free validation over the prescribed COD-LR matrix. (a) Retrieved LR as a function of prescribed LR; the dashed line denotes agreement and marker size increases with prescribed COD. (b) Mean LR bias as a function of prescribed COD.

thetic experiment therefore shows that successful detection of a layer does not imply that its LR is equally well constrained: layer identification remains comparatively stable even when the uncertainty of the optical inversion becomes large.




Figure 2: Propagation of measurement noise through the retrieval chain. (a) Truth-matched detection and boundary-localization success, (b) cloud-base and cloud-top RMSE, and (c) distribution of the DATA_ONLY LR error. The solid curve in panel (c) denotes the median retrieval error and the shaded region spans the 16th–84th percentiles; the dashed horizontal line indicates zero LR error.

The aggregate noise-dependent performance is summarized in Table 5. The transmission-based COD retrieval degrades progressively with decreasing SNR, with RMSE increasing from approximately 5×10^{-6} in the deterministic experiment to 0.2058 at SNR 2. The corresponding DATA_ONLY LR RMSE increases more rapidly, from 5.54 to 42.29 sr, confirming that LR is substantially more sensitive to loss of information in the elastic signal.

The static-prior experiment demonstrates how regularization modifies this noise sensitivity. Under strong data constraint, the broad synthetic prior has little influence on the retrieved LR. As the SNR decreases, however, the prior increasingly suppresses large noise-driven excursions and reduces the LR RMSE, reaching 14.61 sr at SNR 2. This improvement is accompanied by shrinkage toward the prescribed prior centre and must therefore be interpreted as enhanced stability rather than as additional information contained in the elastic measurement itself.

The temporal experiments provide a complementary assessment of regularization by exam-

Table 5: Static synthetic-validation results as a function of target SNR. The PRIOR column corresponds to the controlled broad-prior sensitivity experiment and is not intended to represent a climatological cirrus prior.

SNR	Matched (%)	Boundary ± 150 m (%)	COD _{trans} RMSE	DATA LR RMSE (sr)	PRIOR LR RMSE (sr)
∞	100.0	100.0	0.000005	5.54	5.54
50	100.0	100.0	0.0116	9.65	8.85
20	100.0	99.9	0.0295	13.66	10.55
10	100.0	99.0	0.0527	18.80	11.63
5	100.0	98.6	0.0998	28.32	13.65
2	99.4	98.3	0.2058	42.29	14.61

ining whether continuity between successive retrievals improves the LR estimate. When the prescribed LR remains constant or varies gradually, the temporal constraint suppresses profile-to-profile variability and can improve retrieval stability. Its benefit is not universal, however. Following an abrupt LR transition, the previous retrieval no longer represents the current state and a temporal-only constraint can delay the adjustment of the solution. The step experiment therefore illustrates the principal limitation of temporal regularization: its effectiveness depends on the relationship between the atmospheric evolution time scale and the measurement cadence.

Under stronger noise, combining static and temporal constraints provides the most stable retrieval across the tested temporal sequences, because the measurement itself carries less information and the auxiliary constraints play a larger role. The complete paired static-prior distributions, prescribed temporal sequences, and configuration-by-configuration RMSE comparisons are reported in Appendix A. The main text therefore retains the principal validation results and their interpretation, while the detailed sensitivity analysis is presented separately.

A Synthetic atmosphere and forward-model details

The synthetic validation is designed as a forward–inverse experiment in which the atmospheric state is prescribed independently of the retrieval. Cloud geometry, cloud optical depth (COD), and lidar ratio (LR) are first defined as known truth quantities and used to generate an elastic-lidar signal. The resulting profile is then processed by the same LACRIT chain described in sec.3 without providing any of the prescribed cloud properties to the retrieval. Synthetic truth is used only after processing to evaluate detection and optical-property errors.

A.1 Prescribed molecular and cloud optical state

The synthetic atmosphere consists of a molecular background onto which a single cirrus-like layer is superimposed. Molecular extinction $\alpha_m(z)$ and backscatter $\beta_m(z)$ are calculated from a standard-atmosphere pressure and temperature profile at 532 nm. The corresponding attenuated molecular return is

$$\beta_{m,att}(z) = \beta_m(z) \exp \left[-2 \int_0^z \alpha_m(r) dr \right], \quad (13)$$

which provides the molecular reference subsequently used by LACRIT for normalization, molecular-consistency testing, and transmission retrieval.

The cloud is introduced independently of the molecular atmosphere so that COD and LR remain controlled truth variables. Within the prescribed layer boundaries z_b and z_t , altitude is expressed through the normalized coordinate

$$x(z) = \frac{z - z_b}{z_t - z_b}, \quad (14)$$

and the vertical distribution of particulate extinction is represented by the smooth weighting function

$$w(z) = \begin{cases} \sin^2(\pi x(z)), & z_b \leq z \leq z_t, \\ 0, & \text{otherwise.} \end{cases} \quad (15)$$

Rather than prescribing the peak extinction, the profile is normalized so that its vertical integral reproduces exactly the prescribed cloud optical depth,

$$\alpha_p(z) = \text{COD}_{\text{true}} \frac{w(z)}{\int w(z') dz'}. \quad (16)$$

The particulate backscatter coefficient is then determined independently from the prescribed lidar ratio as follows

$$\beta_p(z) = \frac{\alpha_p(z)}{\text{LR}_{\text{true}}}. \quad (17)$$

This construction is important for the validation because COD_{true} controls the integrated particulate extinction, whereas LR_{true} independently determines the corresponding backscatter. LACRIT must therefore recover both quantities from the simulated elastic signal rather than from parameters introduced directly into the inverse calculation.

The complete atmospheric optical state is finally obtained as

$$\alpha(z) = \alpha_m(z) + \alpha_p(z), \quad \beta(z) = \beta_m(z) + \beta_p(z). \quad (18)$$

A.2 Generation of the synthetic elastic-lidar signal

Once the molecular and particulate radiative properties have been prescribed, the corresponding lidar observation is generated using the single-scattering elastic-lidar equation. The clean range-corrected signal is written as

$$S_{\text{clean}}(z) = K O(z) [\beta_{\text{m}}(z) + \beta_{\text{p}}(z)] \exp \left[-2 \int_0^z (\alpha_{\text{m}}(r) + \alpha_{\text{p}}(r)) dr \right], \quad (19)$$

where K represents the system constant and $O(z)$ the overlap function. Because the synthetic cloud is located near 10 km, full overlap is assumed throughout the altitude interval relevant to the validation, such that $O(z) = 1$. The multiplicative factor K is subsequently removed by the normalization applied by LACRIT and therefore does not affect the retrieval.

Equation 19 contains both effects through which the cloud modifies the elastic return: the increase in backscatter produced within the layer and the two-way attenuation accumulated along the propagation path. The resulting profile consequently provides the observable from which LACRIT must identify the layer boundaries and recover its integrated optical properties.

A.3 Illustration of the detection procedure

Once generated, the synthetic profiles are processed without using the prescribed cloud boundaries. Figure 3 illustrates how the detection criteria defined in sec. 4 act on a representative profile under degraded signal conditions. The example is intended to connect the formal detection criteria to their effect on an actual lidar profile rather than to introduce additional algorithmic rules.

The procedure first identifies a statistically significant cloud core using the strict SNIP-residual and molecular-reference thresholds. The accepted core is subsequently expanded using the lower growth criteria, while the persistence condition prevents isolated bins below threshold from terminating the layer prematurely. Minimum-thickness criteria are then applied to the initial core and to the final expanded layer, and sufficiently close expanded regions are merged to avoid artificial fragmentation of a single cloud structure.




Figure 3: Illustration of the LACRIT cloud-detection criteria using a representative synthetic profile under degraded signal conditions. Panel (A) shows the two-stage hysteresis procedure applied to the SNIP residual, with a restrictive criterion for cloud-core identification and a lower threshold for boundary growth. Panels (B1)–(B2) compare the measured profile with the SNIP and attenuated molecular references and illustrate the independent molecular-consistency criterion. Panels (C1)–(C2) show the minimum core- and final-layer thickness requirements, while panel (D) illustrates the merging of nearby expanded regions.

B Monte Carlo implementation and truth metrics

The deterministic experiment establishes the intrinsic behavior of the retrieval in the absence of measurement noise. To evaluate how this behavior evolves as the observational constraint degrades, the clean synthetic profiles are subsequently perturbed with controlled additive Gaussian noise,

$$S_{\text{noisy}}^{(n)}(z) = S_{\text{clean}}(z) + \varepsilon^{(n)}(z), \quad \varepsilon^{(n)} \sim \mathcal{N}(0, \sigma_n^2), \quad (20)$$

where the superscript n denotes an independent Monte Carlo realization. The noise amplitude is defined relative to the clean molecular signal in the 7–8 km normalization region. More specifically, if S_{ref} denotes the median molecular signal within this interval, the imposed noise standard deviation is

$$\sigma_n = \frac{S_{\text{ref}}}{\text{SNR}_{\text{target}}}. \quad (21)$$

This definition provides a common reference for controlling the perturbation level across the validation matrix while allowing the effective local SNR to vary naturally with altitude and signal magnitude. The noise standard deviation is therefore kept constant along a given synthetic profile rather than adjusted independently at each range gate. The noise-free case is evaluated once for each prescribed COD–LR combination, whereas 50 independent realizations are generated for every finite target SNR and truth case.

To isolate the effect of regularization from stochastic variability, the same noise realization is used when a given synthetic atmosphere is processed with different retrieval configurations. DATA_ONLY, PRIOR_ONLY, TEMPORAL_ONLY, and PRIOR_TEMPORAL can therefore be compared on a paired basis, such that differences between configurations arise from the imposed constraints rather than from differences in the simulated measurement noise.

Performance evaluation is conducted only after the complete LACRIT processing chain has been applied. In particular, the prescribed cloud boundaries are not used during detection or optical retrieval. Retrieved and prescribed layers are associated a posteriori through the one-dimensional intersection-over-union metric,

$$\text{IoU} = \frac{|L_{\text{ret}} \cap L_{\text{true}}|}{|L_{\text{ret}} \cup L_{\text{true}}|}, \quad (22)$$

where L_{ret} and L_{true} denote the retrieved and prescribed vertical layer intervals, respectively. A retrieval is considered truth matched when $\text{IoU} \geq 0.50$. This criterion establishes whether the detected feature corresponds to the prescribed cloud, while the accuracy of its geometrical localization is evaluated independently by requiring both the retrieved cloud base and cloud top to lie within 150 m of their prescribed positions.

Once a layer has been successfully matched, the accuracy of the retrieved radiatives quantities is quantified relative to the prescribed truth. For a retrieved scalar quantity x_i and its corresponding truth value x_i^* , the retrieval error is defined as

$$e_i = x_i - x_i^*, \quad (23)$$

and the root-mean-square error is calculated as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N e_i^2}. \quad (24)$$

RMSE is complemented by the mean bias and by the 16th–84th percentile interval of the individual retrieval errors in order to characterize both systematic displacement and dispersion of the solution. Detection success, boundary accuracy, unmatched candidate layers, and the

fraction of valid optical retrievals are reported separately. This distinction is important because successful identification of the prescribed cloud does not necessarily imply that its COD or LR is equally well constrained. The validation metrics therefore separate geometrical detection performance from the accuracy and stability of the subsequent optical-property retrieval.

C Regularization sensitivity experiments

The influence of regularization is evaluated independently of the baseline retrieval in order to distinguish the information provided by the measurements from that introduced through additional constraints. Two complementary experiments are considered. The first examines the effect of a static LR prior on otherwise independent noisy profiles, while the second evaluates temporal regularization using explicitly ordered LR sequences.

The static-prior experiment is performed on the same noisy COD–LR matrix used for the baseline validation. Its purpose is not to reproduce a specific climatological prior, but rather to isolate the effect of introducing a soft constraint on LR under controlled conditions. A single prior is therefore applied to all truth cases, with $L_0 = 45$ sr and a deliberately broad width of $\sigma_L = 25$ sr. Because this prior is imposed independently of the prescribed LR truth, any improvement in retrieval stability can be attributed to the regularizing effect of the constraint rather than to prior knowledge tuned to the simulated case.

The prior penalty is formulated in logarithmic LR space so that deviations from the prior centre are expressed in relative rather than absolute terms. The linear-space width σ_L is therefore converted to its logarithmic equivalent according to

$$\sigma_{\ln L} = \sqrt{\ln \left[1 + \left(\frac{\sigma_L}{L_0} \right)^2 \right]}. \quad (25)$$

DATA_ONLY and PRIOR_ONLY are evaluated on identical noisy profiles. This paired design allows the effect of the static prior to be quantified directly, including both the reduction in retrieval variance and the associated shrinkage of the solution toward L_0 .

Temporal regularization is investigated separately because the static COD–LR matrix contains independent atmospheric cases and therefore does not define a meaningful temporal evolution. Instead, three ordered LR sequences are constructed at fixed COD = 0.10, using the same 10.0–10.6 km cloud geometry, a 120-s cadence, and 12 profiles per sequence. The scenarios are designed to probe three distinct behaviors: suppression of retrieval jitter under constant conditions, tracking of gradual LR evolution, and the tendency of temporal smoothing to lag behind an abrupt change.

Table 6: Prescribed LR sequences used to evaluate temporal regularization.

Scenario	Prescribed LR sequence (sr)	Purpose
Constant	30, 30, 30, 30, 30, 30, 30, 30, 30, 30, 30, 30, 30	Jitter suppression
Gradual	20, 20, 25, 25, 30, 30, 35, 35, 40, 40, 45, 45	Slow tracking
Step	20, 20, 20, 20, 20, 20, 50, 50, 50, 50, 50, 50	Abrupt-change lag

The temporal constraint is propagated only through the previous *retrieved* LR obtained for the same regularization mode, SNR, and Monte Carlo realization. The prescribed synthetic truth is never supplied to the temporal term, and the first profile in each sequence is processed without temporal regularization. The four configurations DATA_ONLY, PRIOR_ONLY, TEMPORAL_ONLY, and PRIOR_TEMPORAL therefore differ only in the inclusion of the static-prior and temporal penalties, while the underlying atmospheric state and noise realization remain identical.

This design makes it possible to distinguish three effects: the information contained in the lidar profile itself, the stabilizing influence of an external static constraint, and the additional assumption of temporal continuity. The temporal experiment therefore tests not only whether regularization reduces retrieval noise, but also whether that stabilization is obtained at the cost of delayed response when the true LR evolves rapidly.

D Regularization sensitivity results

The static-prior experiment illustrates the expected bias–variance trade-off introduced by regularization. In the deterministic limit, the retrieval is already strongly constrained by the synthetic measurement and the prior has little effect on the recovered LR. As the SNR decreases, however, the measurement constraint becomes progressively weaker and the influence of the prior becomes more apparent. Figure 4a shows that this additional constraint substantially reduces LR RMSE under noisy conditions.




Figure 4: Static-prior sensitivity. (a) LR RMSE for DATA_ONLY and PRIOR_ONLY. (b) Median individual LR error and 16th–84th percentile interval for the same paired Monte Carlo realizations. The dashed horizontal line denotes zero error.

The mechanism behind this improvement is shown in Fig. 4b. Without regularization, the distribution of individual LR errors broadens rapidly as the signal degrades. Introducing the static prior limits this dispersion but simultaneously shifts the retrieved solutions toward the imposed prior centre. The resulting reduction in RMSE should therefore be interpreted as an increase in retrieval stability rather than as the addition of independent LR information. The experiment demonstrates how a soft external constraint can suppress noise-driven variability, while making explicit the associated risk of shrinkage when the prior differs from the true atmospheric state.

The temporal experiments show that continuity is most beneficial when the atmospheric state evolves slowly relative to the retrieval cadence. Under nearly stationary conditions, temporal regularization suppresses profile-to-profile fluctuations effectively. At SNR 10, TEMPORAL_ONLY provides the lowest RMSE for the constant-LR sequence (6.19 sr; Table 7), compared with 11.53 sr for DATA_ONLY. For the gradual sequence, the combined PRIOR_TEMPORAL configuration performs best (5.89 sr), indicating that the temporal constraint can stabilize the retrieval while still allowing a slowly varying LR to be followed.

Table 7: LR RMSE (sr) for the temporal regularization experiments. The lowest RMSE for each sequence and SNR is shown in bold.

SNR	Sequence	DATA_ONLY	PRIOR_ONLY	TEMPORAL_ONLY	PRIOR_TEMPORAL
10	Constant	11.53	9.32	6.19	7.24
	Gradual	12.14	9.02	10.74	5.89
	Step	12.44	9.23	17.14	10.55
5	Constant	24.76	13.81	13.35	11.28
	Gradual	25.30	13.21	16.86	9.43
	Step	26.46	13.93	21.33	13.03




Figure 5: LR RMSE for constant, gradual, and step LR sequences at (a) SNR 10 and (b) SNR 5. The grouped bars show that the relative benefit of each regularization configuration depends jointly on measurement quality and on the temporal evolution of the prescribed LR.

The step experiment illustrates the corresponding limitation. At SNR 10, PRIOR_ONLY yields the lowest RMSE (9.23 sr; Table 7), whereas configurations containing the temporal term respond less efficiently to the abrupt change because the continuity constraint retains information from the preceding state. When the signal is further degraded to SNR 5, the combined PRIOR_TEMPORAL configuration gives the lowest RMSE for all three sequences (Table 7). Its benefit nevertheless remains smaller for the step sequence than for the gradual case, reflecting the trade-off between temporal stabilization and responsiveness to rapid changes.

Overall, temporal regularization should therefore be interpreted as a stabilizing assumption rather than as an additional LR observation. Its benefit depends jointly on measurement quality and on the relationship between the atmospheric evolution time scale and the retrieval cadence. The complete comparison is summarized numerically in Table 7 and graphically in Fig. 5.

E Scope and limitations of the synthetic validation

The synthetic validation is intentionally designed to isolate the behavior of the retrieval under controlled conditions rather than to reproduce the full complexity of the CE376 measurement chain. The resulting experiment therefore provides a truth-based assessment of algorithmic performance within a well-defined set of assumptions, while several physical and instrumental effects remain outside its scope.

The first limitation concerns the representation of radiative transfer. The forward model assumes single scattering and therefore excludes multiple-scattering effects that may become

relevant in optically thicker cirrus. Any residual LR bias observed in the deterministic high-COD cases must consequently originate from the behavior of the retrieval under the adopted single-scattering formulation and cannot be interpreted as a multiple-scattering correction. The synthetic experiment is therefore suited to identifying intrinsic numerical or methodological biases, but not to quantifying radiative-transfer effects that are absent from the forward model itself.

A second simplification concerns the representation of measurement noise. The Monte Carlo perturbation is additive and Gaussian, with an altitude-independent standard deviation referenced to the molecular signal in the normalization region. This approach provides direct control of the effective noise level and allows the propagation of signal degradation through the complete retrieval chain to be evaluated systematically. It does not, however, reproduce the complete photon-counting statistics, detector background, dead-time effects, or channel-specific electronic response of the instrument. The resulting SNR experiments should therefore be interpreted as controlled sensitivity tests rather than as a complete detector-level noise simulation.

The prescribed cloud morphology is likewise simplified. The static validation matrix uses a single smooth 600-m vertical structure while varying COD, LR, and SNR independently. This design isolates the dependence of the retrieval on optical depth, lidar ratio, and signal quality, but it does not span the full range of vertical heterogeneity, multilayer structures, sharp internal gradients, or partially resolved cloud boundaries encountered in observed cirrus. The reported performance therefore characterizes robustness over the controlled COD–LR–SNR space rather than over the complete morphological diversity of natural clouds.

The geometrical criteria used by the automatic detector must also be interpreted as algorithmic settings rather than as physical definitions of cirrus. The 45-m minimum core thickness, 150-m minimum final-layer thickness, four-bin persistence criterion, and 90-m merging distance are introduced to stabilize the automatic identification of coherent vertical structures. Their role is to reduce fragmentation and noise-driven detections, not to impose a universal physical scale on cirrus clouds.

Similarly, the attenuated-molecular consistency criterion improves the rejection of slowly varying non-molecular contributions by requiring detected signal enhancements to exceed the expected molecular return. This test does not provide an unambiguous aerosol–cloud discrimination, since sufficiently strong particulate layers may also exceed the molecular reference. The selection of representative molecular normalization regions therefore remains important when the method is applied to real atmospheric profiles.

Despite these simplifications, the validation preserves a strict separation between prescribed truth and retrieved information. Synthetic truth is used only after the complete LACRIT processing chain has been applied and is never supplied to normalization, SNIP baseline construction, the molecular consistency test, cloud-boundary detection, transmission COD retrieval, PI baseline estimation, or LR optimization. Likewise, the static-prior parameters are imposed independently of the prescribed LR, and temporal regularization propagates the previous retrieved LR rather than the previous true LR.

This separation is essential to the interpretation of the experiment. The agreement obtained between retrieved and prescribed quantities therefore constitutes a truth-based validation of the processing chain under the stated synthetic assumptions, rather than a self-consistency test in which the retrieval is constrained by information derived from the prescribed solution.

References

- Cadet, B., Giraud, V., Haeffelin, M., Keckhut, P., Rechou, A., and Baldy, S.: Improved retrievals of the optical properties of cirrus clouds by a combination of lidar methods, *Appl. Opt.*, 44, 1726–1734, <https://doi.org/10.1364/AO.44.001726>, 2005.