

Author comments on the manuscript “Sentinel-5 SWIR Detector Persistence on Trace Gas Retrievals”, reviewer 1

We thank the referee for the thoughtful assessment of our work and for highlighting considerations regarding its scope and positioning. We appreciate the recognition of the technical quality and clarity of the manuscript. While the study is focused on Sentinel-5, it provides a systematic quantification of detector persistence effects and their impact on trace-gas retrievals, which we believe is of broader relevance to the remote sensing community. We have nevertheless carefully considered the referee’s suggestions and addressed the minor comments below to further improve the manuscript. Page, line, and figure numbers refer to the original manuscript unless stated otherwise.

General Comments:

1. **Lines 13-14, and elsewhere: The abstract and conclusions refer to CH₄ errors as those from the full-physics retrieval and just mention “proxy” without reference to what it means, but a strength of Sentinel-5 is precisely that one can use the proxy method for CH₄ (unlike TROPOMI). The proxy product will receive far more use than the full-physics product for CH₄ and it should be presented as the default product for CH₄.**

Adjusted

According to the Sentinel-5 CH₄ Algorithm Theoretical Basis Document (ATBD) from ESA [Landgraf et al. 2025], both the full-physics and proxy retrieval approaches are described and foreseen as part of the Level-2 CH₄ product, without one being formally defined as the default. In this context, we treat both approaches on an equal footing in our analysis and discussion. That said, we agree with the referee that the term “proxy” was not sufficiently clear. We have therefore revised the manuscript to improve the clarity of its description in Section 2.3, as well as the abstract and conclusions.

We changed the sentence in P1, L8 from:

“This study quantifies persistence-induced biases for CH₄, CO, and the CH₄/CO₂ proxy ratio using the operational RemoTeC CH₄ retrieval algorithm, and also examines persistence errors for a potential CO₂ product.”

to:

“This study quantifies persistence-induced biases for CH₄, CO, and a potential CO₂ product from full-physics retrievals, as well as for the CH₄/CO₂ proxy ratio, whose relative error is representative of proxy CH₄ retrievals, using the operational RemoTeC CH₄ retrieval algorithm.”

Changed P1, L13 from:

“0.71 % for the proxy “

to:

“0.71 % for the CH₄/CO₂ proxy ratio “

Added P6, L175:

“The persistence error affecting the proxy product originates from errors in the retrieved CH₄ and CO₂ columns entering Eq. (1). Since the proxy XCH₄ is directly proportional to the ratio $[CH_4]/[CO_2]$, the relative error in this ratio propagates one-to-one into the relative error of the proxy CH₄ product. Therefore, in the following, we quantify persistence effects for the proxy retrieval in terms of the relative error in the CH₄/CO₂ ratio.”

Changed P7, L199:

“and 1 pixel for the proxy ratio.”

To:

“and 1 pixel for the CH₄/CO₂ proxy ratio.”

Changed P9, L271 (also included changes from Comment 2):

“This study quantified the impact of detector persistence on Sentinel-5 CH₄, CO₂, CO, and CH₄/CO₂ proxy retrievals under realistic surface conditions. The persistence effect—caused by charge trapping and de-trapping charges in the SWIR detector— was simulated to assess its propagation into Level-2

trace-gas products. Persistence-induced biases are most pronounced over low-albedo, high-contrast scenes such as lakes, coastal zones, or other mixed water–land pixels. Although mean biases remain small, the maximum localized deviations—1.59 % for CH₄, 4 % for CO, 1.55 % for CO₂ and 0.71 % for the proxy—constitute substantial fractions of the CO and proxy uncertainty budget (20 %, 1 %) and exceed those for CH₄ and CO₂ (1 %, 0.5 %). These results highlight the need for effective mitigation in operational processing.”

To:

“This study quantified the impact of detector persistence on Sentinel-5 CH₄, CO₂, and CO full-physics retrievals, as well as on proxy CH₄ retrievals based on the CH₄/CO₂ ratio, under realistic surface conditions. The persistence effect—caused by charge trapping and de-trapping in the SWIR detector—was simulated to assess its propagation into Level-2 trace-gas products. Persistence-induced biases are most pronounced in high-contrast scenes, such as lakes, coastal zones, and mixed land–water pixels. Overall, the biases remain small at scene level, with typical amplitudes below 0.12 % for CH₄, 0.10 % for CO₂, 0.34 % for CO, and 0.06 % for the CH₄/CO₂ proxy ratio. Larger deviations occur in spatially structured regions, where localized errors can reach up to 1.59 % for CH₄, 1.55 % for CO₂, 4 % for CO, and 0.71 % for the CH₄/CO₂ proxy ratio. While these effects represent a noticeable contribution to the error budget in such conditions—particularly for CO and, to a lesser extent, for the proxy—they do not dominate the overall error budget under typical observing conditions. Nevertheless, their spatial correlation with surface contrast may introduce patterns that could be misinterpreted as localized enhancements, highlighting the importance of targeted mitigation strategies such as filtering of affected high-contrast pixels in operational processing.”

2. **Lines 13-14, and conclusions: Giving maximum errors in abstract/conclusions is misleading without the error statistics, which would show that the errors are in general considerably smaller than these maxima. In fact, what I get from this work is that the detector lag will not significantly affect my use of the proxy product for CH₄ – I have other errors to worry about.**

Adjusted

We agree that quoting only maximum errors in the abstract and conclusions can overemphasize worst-case situations. We have therefore revised the text to better distinguish between typical errors and the larger deviations found in specific high-contrast scenes. In particular, we now emphasize that the scene-wide bias statistics are generally small, while the largest errors are concentrated in mixed and high-contrast pixels such as coastal and land–water interface regions. We have also moderated the wording regarding the impact on the proxy product to reflect that persistence is not the dominant error source under most conditions, although it remains relevant in high-contrast scenes, where persistence-induced errors are spatially structured and correlated with surface features, potentially introducing patterns that could be misinterpreted as localized CH₄ enhancements depending on the application.

We changed in P1, L11:

“Across these scenes, the induced biases remain small, with typical amplitudes below 0.12 % for CH₄, 0.10 % for CO₂, 0.34 % for CO, and 0.06 % for the CH₄/CO₂ proxy. However, localized outliers reaching up to 1.59 % for CH₄, 1.55 % for CO₂, 4.01 % for CO and 0.71 % for the proxy constitute a substantial fraction of the CO and the proxy performance targets, and exceed those defined for CH₄ and CO₂. Our findings suggest that targeted post-filtering, such as omitting coastal and land–water interface pixels, can effectively reduce the most significant outliers in the persistence bias, thus aligning the retrieval biases with the required mission performance. “

To:

“Across these scenes, persistence-induced biases are generally small, with typical amplitudes below 0.12 % for CH₄, 0.10 % for CO₂, 0.34 % for CO, and 0.06 % for the CH₄/CO₂ proxy. However, larger deviations occur in high-contrast scenes, particularly at coasts and land–water transitions, where localized errors can reach up to 1.59 % for CH₄, 1.55 % for CO₂, 4.01 % for CO, and 0.71 % for the proxy. These values represent a substantial fraction of the performance targets for CO and the proxy, and exceed those for CH₄ and CO₂. This indicates that, while persistence is not the dominant error source under most conditions, it can still produce non-negligible biases in spatially structured scenes, which may be misinterpreted as localized enhancements. Therefore, such effects should be mitigated, for example through targeted quality control or filtering of affected pixels.”

Changed in P9,L271:

This change affects the same paragraph as discussed above in Comment 1. To avoid repetition, we refer to the revised version provided in the previous response, which incorporates all changes.

3. **Figure 1 is not very helpful. It would be more helpful without unfamiliar abbreviations and with a more substantial caption that describes the overall OSSE strategy.**

Adjusted

We agree that the original version of Figure 1 was not sufficiently self-explanatory. The figure has been revised to improve clarity by replacing abbreviations (e.g. L1B and L2). We have also expanded the caption to better describe the overall OSSE strategy illustrated in the figure.

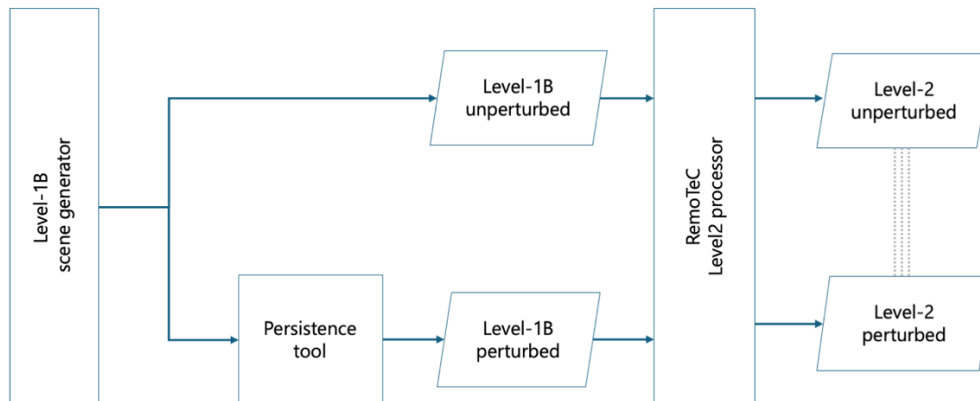


Figure 1. Schematic overview of the simulation framework used to assess the impact of persistence perturbations. Synthetic spectra are generated from a common reference scene and propagated through two processing branches: one including persistence perturbations and one without perturbations. The resulting Level-1B radiance spectra are processed independently through the retrieval chain to produce Level-2 products. The comparison between the perturbed and unperturbed Level-2 outputs enables quantification of the impact of persistence effects on the retrieved quantities.

4. **Line 120: I didn't understand $\text{reff} = 4$. Is it missing a unit? But μm would be too large and nm too small?**

Adjusted

In the original manuscript, the parameter was incorrectly described as an effective radius. In our setup, the aerosol size distribution is prescribed as a power-law distribution following Butz et al. (2010), and the parameter value refers to the exponent of this distribution rather than to a particle radius. It is therefore dimensionless. We have revised the text accordingly to correct this description and avoid confusion.

We have changed P.4, L.118 from:

“The synthetic spectra include a mild aerosol load with an aerosol optical thickness of 0.1 at 760 nm, represented by a Gaussian vertical distribution centered near the surface (Butz et al., 2010), and assuming standard size distribution with $\text{reff} = 4$ (Butz et al., 2011) and optical properties taken from lookup tables (Dubovik et al., 2006).”

To:

“The synthetic spectra include a mild aerosol load with an aerosol optical thickness of 0.1 at 760 nm, represented by a Gaussian vertical distribution centered near the surface (Butz et al., 2010), with aerosol microphysical properties described by a power-law particle size distribution ($\alpha = 4$; Butz et al., 2011) and optical properties taken from lookup tables (Dubovik et al., 2006).”

5. **Line 120: I didn't understand “Surface albedo spectral dependence is modeled as wavelength-independent.” Does this mean that the albedo spectra are just scaled up or down depending on surface? That wouldn't make sense because the wavelength dependence of the albedo should vary with surface type.**

Adjusted

We agree that, in reality, the spectral dependence of surface albedo varies with surface type. In our study, however, the surface albedo is treated as a constant value within each spectral band, while varying between bands and from pixel to pixel. This simplification is applied consistently in both the forward simulations and the retrieval. This is motivated by the fact that the Sentinel-2 data used to construct the reference scenes provide albedo values integrated over broad spectral bands, rather than spectrally resolved reflectance. As a result, no wavelength-dependent albedo within each band can be derived in a consistent way. The albedo is therefore treated as spectrally constant within each band in the forward simulations, and the same assumption is adopted in the retrieval to ensure that any differences can be attributed solely to the persistence effect. We have revised the text to clarify this point.

We changed line in P.4, L120 from:

“Surface albedo spectral dependence is modeled as wavelength-independent.”

To:

“Surface albedo is treated as constant within each spectral band (i.e. wavelength-independent), while varying between bands and from pixel to pixel, consistent with the band-averaged reflectance provided by Sentinel-2.”

Added on P7, L185 the following line, where we have also addressed comment from referee 1:

“The same albedo parameterization as in the forward simulations is adopted in the retrieval. For trace gases and, where applicable, aerosols, the a priori is set equal to the true state used to generate the synthetic spectra, and no explicit a priori uncertainties are imposed. This configuration provides a consistent reference case, so that the dominant differences between perturbed and unperturbed retrievals can be attributed to the persistence effect.”

- 6. Line 199 and elsewhere. I don't understand why the detector lag would affect the CH₄ proxy retrieval at all, since the albedo cancels between CO₂ and CH₄, except for some smearing which is maybe what the “one pixel” refers to? But that effect would be independent of change in albedo and it would not translate into a systematic bias.**

Adjusted

We agree that the proxy method reduces errors that are common to both CH₄ and CO₂ retrievals. However, the cancellation is not complete. In our case, detector persistence affects the measured radiances in a scene- and wavelength-dependent manner. Since CH₄ and CO₂ are retrieved from different spectral regions with different absorption features and radiance levels, the persistence-induced perturbations are not identical for the two gases. As a result, they do not fully cancel in the proxy ratio and can lead to residual biases. We have clarified this point in the manuscript.

We added on P7, L207 the following text:

“The non-zero bias in the proxy ratio further indicates that persistence-induced perturbations differ between the spectral regions used for CH₄ and CO₂ retrievals, and therefore do not fully cancel in the ratio.”

And on P9, L276:

“The presence of residual biases in the proxy product indicates that persistence effects do not fully cancel between the CH₄ and CO₂ retrievals.”

- 7. Line 205: I don't understand the nits unit. Aren't the % errors reported for column mixing ratio?**

Adjusted

The unit “nits%” was a typographical error. The proxy bias is expressed as a relative error, consistent with the other quantities, and has been corrected to “%” in the revised manuscript.

8. **Line 260 and elsewhere: I think that the authors' recommendation to exclude fractional water scenes is too radical and would result in unnecessary loss of information for important locations such as coastal cities. The errors reported are not that large, and I certainly would not want to filter out proxy data on that basis.**

We agree that a strict exclusion of fractional water scenes may not be desirable in all applications, particularly for regions of interest such as coastal areas. Our intention was not to propose a rigid filtering criterion, but rather to illustrate a possible mitigation strategy to reduce the largest persistence-induced biases. In practice, such filtering can be implemented through quality flags, allowing users to balance data coverage and accuracy depending on their specific application. We have revised the text to clarify that this approach is not prescriptive and that more refined filtering strategies can be developed.

In this context, it is worth noting that filtering of mixed land–water pixels is already commonly applied in operational retrievals (e.g. via land-fraction-based quality flags in the CH₄ TROPOMI product; SRON, 2025), albeit for different reasons.

We changed on P10, L286:

“This demonstrates that such masking can substantially limit persistence-induced artifacts in the operational retrievals, providing a practical short-term solution.”

To:

“While such masking is effective in limiting the largest persistence-induced artifacts, it may also reduce data coverage in regions of interest such as coastal areas. In practice, this trade-off can be managed through quality indicators, allowing users to balance data coverage and accuracy depending on their application.”

References

Landgraf, J., Butz, A., Hasekamp, O., Hu, H., aan de Brugh, J., Sentinel 5 L2 Prototype Processors, Algorithm Theoretical Baseline Document: Methane Retrieval, SRON-ESA-S5L2PP-ATBD-001-v3.1-20190517-CH₄, SRON Netherlands Institute for Space Research, Leiden, The Netherlands, 2025

Butz, A., ., et al. (2010), CH₄ retrievals from space-based solar backscatter measurements: Performance evaluation against simulated aerosol and cirrus loaded scenes, *Journal of Geophysical Research: Atmospheres*, 115, <https://doi.org/https://doi.org/10.1029/2010JD014514>, 2010.

Butz, A., et al. (2011), Toward accurate CO₂ and CH₄ observations from GOSAT, *Geophys. Res. Lett.*, 38, L14812, doi:10.1029/2011GL047888.

SRON Netherlands Institute for Space Research, 2025. *Sentinel-5P Methane Product Readme File (CH₄ PUM), Version 2.9*. Available at: [https://sentiwiki.copernicus.eu/_attachments/1673595/S5P-MPC-SRON-PRF-CH₄ - Sentinel-5P Methane Product Readme File 2025 - 2.9.pdf](https://sentiwiki.copernicus.eu/_attachments/1673595/S5P-MPC-SRON-PRF-CH4 - Sentinel-5P Methane Product Readme File 2025 - 2.9.pdf)