https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

XDOAS: A Unified Covariance–PCA DOAS Framework for Enhanced Trace Gas Retrievals – Application to OClO from TROPOMI

Leonardo M. A. Alvarado[1], Sora Seo[1], Adrian Doicu[1], Andreas Richter[2], and Diego Loyola[1]

[1]German Aerospace Center (DLR), Remote Sensing Technology Institute, Oberpfaffenhofen, Germany
[2]Institute of Environmental Physics (IUP), University of Bremen, Bremen, Germany

Correspondence: Leonardo M. A. Alvarado (leonardo.alvarado@dlr.de)

Abstract. The treatment of the fit residual is a central problem in DOAS retrievals from hyperspectral satellite instruments. At the precision of instruments such as TROPOMI, the residual of a classical fit is no longer an approximately white background, but splits into systematic structures of physical and instrumental origin and a correlated random noise. Classical DOAS treats both in the same way, which limits its performance for weak absorbers retrieved under difficult conditions. This study presents XDOAS, a complete retrieval that unifies two recent approaches to this problem, covariance weighting and principal component analysis (PCA), within a single linear formulation. The PCA components fit the systematic structures out of the residual explicitly, while a regularised covariance down-weights the remaining correlated noise. Classical DOAS, covariance-weighted DOAS, and PCA-DOAS emerge as limiting cases of XDOAS. The algorithm is applied to the retrieval of OClO slant column densities from TROPOMI/S5P, a weak absorber measured at large solar zenith angles where the residual structures are most pronounced. Evaluated against the three reference retrievals on TROPOMI measurements using identical fit parameters, XDOAS reaches the precision of the best advanced retrievals while being the only one that also gives an uncertainty consistent with the spectral fit quality of each pixel. A consistency test between the RMS residual and the retrieval uncertainty separates the four retrievals, covariance-weighted DOAS gives uncertainties that are decoupled from the residual ($R^2 = 0.002$), PCA-DOAS gives uncertainties that are fixed to the residual because the PCA components fit out part of the random noise ($R^2 = 0.991$), and XDOAS is the only retrieval whose uncertainty follows the spectral fit quality of each pixel with a moderate pixel-to-pixel spread ($R^2 = 0.977$). The comparison against the S5P-PAL OClO product based on empirical post-processed DOAS retrieval, shows close agreement ($R^2 = 0.922$). Where the OClO signal is strong, over the polar vortex, the two products agree almost exactly (slope 1.02, $R^2 = 0.946$), while XDOAS gives a lower RMS residual and a smoother uncertainty field specially, most of all under the low-illumination conditions in which OClO is measured. XDOAS is the baseline algorithm for the operational OClO product of the Copernicus Sentinel-5 mission from the EUMETSAT AC-SAF. XDOAS is a general framework that can be applied to any trace gas with narrow absorption bands, and in particular to measurements under extrema viewing conditions typically in geostationary missions like Sentinel-4, GEMS, and TEMPO.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

1 Introduction

Understanding the variability, depletion, and recovery of stratospheric ozone remains a central topic in atmospheric chemistry. Ozone absorbs harmful ultraviolet radiation and is also involved in a wide range of chemical and radiative processes that affect atmospheric composition and climate. Since the discovery of the Antarctic ozone hole (Farman et al., 1985), long-term observations have been essential for assessing the effectiveness of the Montreal Protocol and for quantifying the influence of changing climate conditions on ozone recovery (Solomon, 1999; World Meteorological Organization, 2022).

Polar ozone loss is mainly driven by catalytic cycles involving reactive halogen species, especially chlorine and bromine radicals, with additional modulation by nitrogen and hydrogen oxide chemistry (Solomon et al., 1986; McElroy et al., 1986). During winter and early spring, the polar vortex isolates cold stratospheric air and favours the formation of polar stratospheric clouds (PSCs) when temperatures fall below approximately 195 K. Heterogeneous reactions on PSC particles convert chlorine reservoir species, mainly HCl and $ClONO_2$, into photolabile compounds. After sunrise, these species release reactive chlorine radicals that initiate efficient catalytic ozone destruction (Solomon et al., 1986; Peter and Grooß, 2012; Höpfner et al., 2018). The magnitude and timing of ozone depletion therefore depend not only on the abundance of inorganic chlorine and bromine, but also on PSC occurrence, denitrification, vortex stability, and interannual dynamical variability. These processes explain much of the difference between the persistent Antarctic ozone hole and the more variable Arctic ozone depletion (Solomon, 1999).

Chlorine dioxide (OClO) is commonly used as an indicator of chlorine activation in polar regions. It is mainly formed through the reaction of ClO with BrO and is rapidly photolysed at low solar zenith angles. Although OClO does not directly cause net ozone destruction, its abundance is closely related to active halogen chemistry under twilight conditions and therefore provides valuable information on chlorine activation inside the polar vortex (Solomon et al., 1987; Schiller and Wahner, 1996; Puķīte et al., 2022). Large OClO columns are generally observed during cold, PSC-active winters and springs, especially in the Southern Hemisphere, while much smaller amounts are usually found outside the polar vortex or during dynamically disturbed Arctic winters. However, comparisons between observations and models have shown that OClO remains difficult to reproduce under conditions of strong chlorine activation and at large solar zenith angles (Oetjen et al., 2011). These limitations highlight the need for robust retrieval methods capable of detecting weak and structured absorption signals under challenging observational conditions.

Differential Optical Absorption Spectroscopy (DOAS) has become one of the most widely used techniques for detecting atmospheric trace gases with narrow absorption structures in the ultraviolet, and visible wavelength ranges. The method separates narrow absorption bands from broad band contributions caused by scattering, aerosols, clouds, and surface reflection as well as slow variation of trace gases. This allows the differential slant column densities (DSCD) of different absorbers to be retrieved from the same measured spectrum (Platt and Stutz, 2008). Early ground-based scattered-light measurements of atmospheric NO_2 were reported by Noxon (1975), while the formal DOAS approach was introduced by Platt et al. (1979) for the simultaneous detection of CH_2O, O_3, and NO_2. Since then, DOAS has evolved from long-path and zenith-sky ground-based observations to multi-axis, airborne, balloon-borne, and satellite-borne applications. The development of MAX-DOAS

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

extended the method by using multiple viewing elevations, increasing sensitivity to lower-tropospheric absorbers and enabling profile retrievals when combined with radiative transfer modelling (Hönninger et al., 2004; Bösch et al., 2018).

Early observations of trace gases from space using the DOAS method were reported by Burrows et al. (1999), based on measurements from the Global Ozone Monitoring Experiment (GOME), which demonstrated that accurate global measurements of ozone (Bramstedt et al., 2003) and several trace gases could be detected from UV–visible backscattered radiation. Later missions, such as the Scanning Imaging Absorption Spectrometer for Atmospheric Chartography (SCIAMACHY), the Global Ozone Monitoring Experiment-2 (GOME-2), the Ozone Monitoring Instrument (OMI), and the TROPOspheric Monitoring Instrument (TROPOMI), opened the way for a variety of observations, leading to improvements in spatial resolution, spectral quality, temporal coverage (Coldewey-Egbers et al., 2022), and the number of retrievable species. In addition, DOAS retrievals have been refined through improvements on wavelength calibration, treatment of the Ring effect, correction of interfering absorbers, wavelength-dependent air-mass-factor formulations, weighting-function approaches, and statistical methods designed to reduce systematic residual structures for several species (Platt and Stutz, 2008; Coldewey-Egbers et al., 2005; Puķīte et al., 2021; Alvarado et al., 2014). More recently, principal component analysis (PCA) has been introduced in satellite retrievals to represent instrumental or scene-dependent spectral variability and to improve sensitivity to weak trace-gas signals (Li et al., 2013; Joiner et al., 2023) as well as a covariance-based spectral retrieval strategies have recently been explored as an alternative or complement to classical DOAS fitting (Theys et al., 2025).

This study presents the eXtended DOAS retrieval for trace gases (XDOAS), a synergistic framework, and demonstrates its performance for the retrieval of OClO, used as the primary case study because it is a weak absorber, is strongly linked to polar chlorine activation, and is often retrieved under challenging conditions, including large solar zenith angles, low signal-to-noise ratios, and strong spectral interference from other absorbers. These characteristics make OClO an ideal target for testing retrieval strategies designed to improve the detection of weak trace-gas signals. The proposed XDOAS framework can be applied to other trace gases with narrow absorption bands, such as CHOCHO, NO_2, HCHO, etc.

Four retrieval approaches are considered: standard DOAS, covariance-DOAS, PCA-DOAS including an adaptive PCA-DOAS configuration, and the proposed covariance–PCA Synergy DOAS method. The manuscript is organized as follows. First, the methodology is described, including the XDOAS framework and its derivation from classical DOAS. Second, the different approaches are compared with conventional DOAS and existing DOAS extensions. Third, the performance of the methods is evaluated using OClO as the reference trace gas. Finally, the retrieval results are compared with the pre-operational TROPOMI OClO product.

2 DOAS retrieval: from classical fit to unified residual correction

The Differential Optical Absorption Spectroscopy (DOAS) method has been widely used in spaceborne observations to retrieve atmospheric trace gases with narrow absorption features in the ultraviolet and visible spectral ranges (e.g., Burrows et al., 1999; Palmer et al., 2001; Oetjen et al., 2011; Richter et al., 2011; Alvarado et al., 2014; González Abad et al., 2015; De Smedt et al., 2018; Alvarado et al., 2020; Puķīte et al., 2021; van Geffen et al., 2022; Puķīte et al., 2022). The method is based on the Beer–

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Lambert law and exploits the fact that many atmospheric absorbers have narrow absorption bands within a selected spectral window (Platt and Stutz, 2008).

The central idea of DOAS is to separate the absorption cross section of each absorber into a slowly varying broadband component and a rapidly varying narrow-band component,

$$\sigma_i(\lambda) = \sigma_i^0(\lambda) + \sigma_i'(\lambda), \tag{1}$$

where $\sigma_i^0(\lambda)$ represents the broadband contribution and $\sigma_i'(\lambda)$ contains the narrow-band absorption structures used for the retrieval. The slowly varying terms (broadband absorption, Rayleigh and Mie scattering, surface reflection, clouds, aerosols, and other smooth instrumental or radiative-transfer effects) are represented by a low-order polynomial. With this separation, the measured intensity ratio can be written as

$$\frac{I(\lambda)}{I_0(\lambda)} = \exp\left[-\sum_{i=1}^{N_g} \sigma_i'(\lambda)\, S_i - \sum_{p=0}^{P} a_p \lambda^p\right], \tag{2}$$

where $I_0(\lambda)$ is the reference spectrum, $I(\lambda)$ the measured radiance, S_i the slant column density (SCD) of absorber i, a_p the polynomial coefficients, and P the polynomial order. Taking the natural logarithm defines the optical depth, and leads to the linear equation

$$\tau^\delta(\lambda) = \ln\left[\frac{I_0(\lambda)}{I(\lambda)}\right] = \sum_{i=1}^{N_g} \sigma_i'(\lambda)\, S_i + \sum_{p=0}^{P} a_p \lambda^p + \delta(\lambda), \tag{3}$$

where $\tau^\delta(\lambda)$ is the measured optical depth and $\delta(\lambda)$ is the *data error*, collecting all contributions to the measured optical depth that are not represented by the trace-gas cross sections or the polynomial: random detector noise, unmodelled instrumental structures, calibration uncertainties, and higher-order radiative-transfer effects not captured by the linear forward model. Only absorbers with structured features within the fitting window can be retrieved directly by DOAS, while smooth spectral variations are absorbed by the polynomial.

Equation (3) is, by construction, linear in the SCDs, S_i, and the polynomial coefficients a_p. In practice this linearity is broken by small wavelength misalignments between the measured and the reference spectra, arising from instrumental thermal drift, Doppler shifts, or imperfect spectral calibration. These are described by a wavelength shift s and a squeeze coefficient q applied to the reference axis, because both appear inside the argument of the differential cross sections, $\sigma_i'(\lambda - s - q(\lambda - \lambda_0))$, they enter the model non-linearly. Following Beirle et al. (2013), for the residual deformations considered here ($|s| \lesssim 5\,\mathrm{pm}$, $|q| \lesssim 10^{-3}$) the shift and squeeze contributions are approximated to first order by a Taylor expansion of the reference spectrum and represented by additional pseudo-absorbers (Sect. 2.1), which enter the fit linearly. The resulting linearised fit is solved in a single matrix inversion, avoiding the iterative refinement and convergence issues of non-linear optimisers, this is particularly relevant for operational retrievals from instruments such as TROPOMI, where the same fit must be performed for millions of pixels per orbit.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

For compact notation, the optical depth and all spectral basis functions are sampled on the detector wavelength grid $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_M)^\top$, yielding the data vector $\boldsymbol{\tau}^\delta \in \mathbb{R}^M$. The linearised fit can then be written in matrix form as

$$\boldsymbol{\tau}^\delta = \mathbf{B}\boldsymbol{\beta} + \boldsymbol{\delta}, \qquad \mathbf{B} \in \mathbb{R}^{M \times N}, \tag{4}$$

where $\mathbf{B}$ contains, as columns, the discretised differential cross sections, the polynomial terms, and any auxiliary pseudo-absorbers included in the fit, and N is the total number of fit parameters. The vector $\boldsymbol{\beta} \in \mathbb{R}^N$ collects all unknown fit coefficients and can be partitioned as

$$\boldsymbol{\beta} = \begin{pmatrix} \mathbf{S} \\ \mathbf{a} \\ \mathbf{c}_{\text{aux}} \end{pmatrix}, \tag{5}$$

where $\mathbf{S} = (S_1, \ldots, S_{N_g})^\top$ contains the SCDs, $\mathbf{a} = (a_0, \ldots, a_P)^\top$ the polynomial coefficients, and $\mathbf{c}_{\text{aux}}$ the coefficients of the auxiliary pseudo-absorbers (shift, squeeze, intensity offset, mean residual, and the PCA components introduced in Sects. 2.4–2.5). The SCDs, $\mathbf{S}$, are the physical quantities of interest, the remaining components of $\boldsymbol{\beta}$ describe broadband, instrumental, and statistical effects that are co-fitted to retrieve $\mathbf{S}$ accurately. The ordinary least-squares (OLS) estimate $\hat{\boldsymbol{\beta}}$, where a hat denotes a fitted (estimated) quantity as opposed to the true state $\boldsymbol{\beta}^\dagger$ introduced below, is the minimiser $\frac{1}{2}\|\mathbf{B}\boldsymbol{\beta} - \boldsymbol{\tau}^\delta\|^2$, giving

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{B}^\top\mathbf{B}\right)^{-1}\mathbf{B}^\top\boldsymbol{\tau}^\delta, \tag{6}$$

from which the retrieved SCDs, $\hat{\mathbf{S}}$, are extracted as the first N_g components of $\hat{\boldsymbol{\beta}}$. The fit residual is the computable quantity

$$\mathbf{r} = \boldsymbol{\tau}^\delta - \mathbf{B}\hat{\boldsymbol{\beta}}. \tag{7}$$

It is important to distinguish the residual $\mathbf{r}$ from the data error $\boldsymbol{\delta}$. Writing the true state vector $\boldsymbol{\beta}^\dagger$, such that $\boldsymbol{\tau}^\delta = \mathbf{B}\boldsymbol{\beta}^\dagger + \boldsymbol{\delta}$, the residual can be expressed as

$$\mathbf{r} = \boldsymbol{\delta} + \mathbf{B}\left(\boldsymbol{\beta}^\dagger - \hat{\boldsymbol{\beta}}\right). \tag{8}$$

The data error is an unknown perturbation affecting the measurements, whereas the residual is obtained after the fit and measures the discrepancy between the measured data and the fitted model. If the preliminary fit is sufficiently accurate, $\hat{\boldsymbol{\beta}} \approx \boldsymbol{\beta}^\dagger$, then $\mathbf{r} \approx \boldsymbol{\delta}$ and the statistics of $\mathbf{r}$ reproduce those of $\boldsymbol{\delta}$. This approximation is what justifies using the empirical covariance of a residual ensemble as a proxy for the unknown data-error covariance in the advanced retrievals introduced below.

2.1 Spectral corrections and wavelength calibration

In practice the DOAS fit applied to satellite measurements includes a set of standard corrections, implemented as additional pseudo-absorbers within $\mathbf{B}$ so as to preserve the linearity of the fit. Before the fit, all reference cross sections are prepared on a common basis: the high-resolution laboratory cross sections are convolved with the instrument spectral response function (ISRF) of TROPOMI, with strict conservation of the slit-function area, and interpolated onto the detector wavelength grid, so

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

that they are directly comparable with the measured spectrum. Following Beirle et al. (2013), the wavelength misalignment is described by a shift s and a stretch/squeeze coefficient q around a reference wavelength λ_0,

$$\lambda \to \lambda - s - q(\lambda - \lambda_0). \tag{9}$$

With this convention $q = 0$ corresponds to a non-deformed axis, $q < 0$ to a local stretch and $q > 0$ to a local squeeze, equivalent to that of Beirle et al. (2013) up to the additive constant $q = 1 - q_{\mathrm{B}}$. Because the misalignment acts on the reference spectrum, it is convenient to write $\tau_{\mathrm{ref}} \equiv \ln I_0$. A first-order Taylor expansion then gives

$$\ln I_0(\lambda - s - q(\lambda - \lambda_0)) \approx \ln I_0(\lambda) - s\frac{\partial \ln I_0}{\partial \lambda} - q(\lambda - \lambda_0)\frac{\partial \ln I_0}{\partial \lambda}, \tag{10}$$

so that the shift and squeeze are represented by the two pseudo-absorbers

$$\mathbf{u}_s = \frac{\partial \ln I_0}{\partial \lambda}, \qquad \mathbf{u}_q = (\lambda - \lambda_0)\frac{\partial \ln I_0}{\partial \lambda}, \tag{11}$$

appended as extra columns of $\mathbf{B}$ and fitted simultaneously with the atmospheric absorbers, the fit then returns s and q directly, without any external non-linear optimisation loop.

Additive radiance offsets (stray light, dark current, residual level-1 effects) are described by a term $I_{\mathrm{off}}(\lambda)$ in the measured radiance, $I_{\mathrm{meas}} = I + I_{\mathrm{off}}$. In the optically thin limit ($I \approx I_0$) the induced perturbation of the optical depth is, to first order, $\Delta\tau \approx -I_{\mathrm{off}}/I_0$, thus, following standard DOAS practice (Theys et al., 2017) the offset is represented by a $1/I_0(\lambda)$ pseudo-absorber appended as an extra column of $\mathbf{B}$.

$$I_{\mathrm{meas}}(\lambda) = I(\lambda) + I_{\mathrm{off}}(\lambda), \qquad \Delta\tau(\lambda) \approx -\frac{I_{\mathrm{off}}(\lambda)}{I_0(\lambda)}. \tag{12}$$

Isolated detector spikes and unflagged bad pixels are removed by iterative spike detection on the fit residual (Richter et al., 2011). After an initial fit, the residual is normalised by its root-mean-square value,

$$r_n^{\mathrm{norm}}(\lambda_k) = \frac{r_n(\lambda_k)}{\sqrt{\frac{1}{N_{\mathrm{dof}}}\sum_k r_n^2(\lambda_k)}}, \tag{13}$$

where $N_{\mathrm{dof}} = M - N$ is the number of degrees of freedom. Points with $|r_n^{\mathrm{norm}}(\lambda_k)| > \kappa$ (typically $\kappa = 5$) are flagged as outliers, replaced by the model value, and the fit is repeated until no further spikes are detected.

2.2 Limitations of the classical residual treatment

The corrections of Sect. 2.1 improve the spectral alignment and mitigate isolated artefacts, but they do not change how the remaining residuals are processed, in the classical DOAS fit all wavelength points are implicitly treated as if the data error had independent components of comparable variance. If the data-error covariance matrix $\mathbf{C}_\delta = \mathbb{E}\left[(\boldsymbol{\delta} - \mathbb{E}[\boldsymbol{\delta}])(\boldsymbol{\delta} - \mathbb{E}[\boldsymbol{\delta}])^\top\right]$ were known (here $\mathbb{E}[\cdot]$ denotes the expectation, i.e. the ensemble average), the best linear unbiased estimator would follow from generalized least squares (GLS),

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{B}^\top \mathbf{C}_\delta^{-1} \mathbf{B}\right)^{-1} \mathbf{B}^\top \mathbf{C}_\delta^{-1} \boldsymbol{\tau}^\delta. \tag{14}$$

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

For modern hyperspectral satellite measurements the white-noise assumption is inadequate. Physical effects (temperature-dependent cross sections, wavelength dependence of the air mass factor, the I_0 effect, higher-order radiative-transfer effects, Puķīte et al., 2010; Rozanov and Rozanov, 2010) and instrumental effects (spectral undersampling, detector ageing, residual wavelength misregistration, ISRF variability, across-track smile, stray light, intensity offsets) generate residual structures correlated over several wavelength pixels. The data error therefore splits into

$$\boldsymbol{\delta} = \boldsymbol{\delta}_{\mathrm{sys}} + \boldsymbol{\delta}_{\mathrm{rnd}}, \qquad \mathbb{E}[\boldsymbol{\delta}_{\mathrm{rnd}}] = \mathbf{0}, \qquad \mathbb{E}[\boldsymbol{\delta}] = \boldsymbol{\delta}_{\mathrm{sys}}, \tag{15}$$

i.e. a small number of reproducible systematic structures of physical or instrumental origin, and a zero-mean correlated random noise. The classical DOAS fit treats both contributions identically, which is its fundamental limitation. The retrieval approaches introduced below address this in complementary ways: covariance-weighted DOAS treats the random part through reweighting (Sect. 2.3), the PCA-based correction handles the systematic structures explicitly (Sect. 2.4), and the unified XDOAS retrieval combines both, treating each contribution accordingly (Sect. 2.5).

Since $\mathbf{C}_\delta$ is not known, it is estimated empirically. A pre-fit is performed over a clean equatorial sector, with the target gas excluded from the fitted absorbers, so that the resulting residual spectra $\mathbf{r}_n$ ($n = 1, \ldots, N_{\mathrm{ref}}$) are free of the target absorption. This construction assumes a region where the target column is expected to be zero or close to zero, as for OClO at low latitudes; for a gas without such a background-free region, the adaptive configuration of Sect. 2.4.1 is used, in which the target absorption is co-fitted and its large-scale structure removed before the covariance and PCA components are derived. The ensemble is built in a first pass over the orbit, and the resulting statistics are applied in a second pass. This follows the CODOAS approach of Theys et al. (2025) rather than the externally referenced covariance of COBRA (Theys et al., 2021). The mean residual and empirical covariance are

$$\bar{\mathbf{r}} = \frac{1}{N_{\mathrm{ref}}} \sum_{n=1}^{N_{\mathrm{ref}}} \mathbf{r}_n, \qquad \mathbf{C} = \frac{1}{N_{\mathrm{ref}} - 1} \sum_{n=1}^{N_{\mathrm{ref}}} (\mathbf{r}_n - \bar{\mathbf{r}})(\mathbf{r}_n - \bar{\mathbf{r}})^\top. \tag{16}$$

Because the mean is subtracted, $\mathbf{C}$ estimates the covariance of the fluctuating residual part. Under the approximation $\mathbf{r} \approx \boldsymbol{\delta}$ and a representative ensemble, $\mathbf{r}_{\mathrm{sys}} \approx \bar{\mathbf{r}}$ and $\mathrm{Cov}(\boldsymbol{\delta}_{\mathrm{rnd}}) \approx \mathbf{C}$.

2.3 Covariance-weighted DOAS

The use of the residual covariance matrix to reweight the spectral fit was pioneered, in the context of DOAS satellite retrievals, by Theys et al. (2021) in the COBRA algorithm and developed further in the CODOAS approach of Theys et al. (2025). The covariance-weighted DOAS considered here follows that line of work. Covariance-weighted DOAS modifies the spectral weighting of the fit using the residual covariance of the reference ensemble. Because the amplitude of the systematic structure in an individual retrieval spectrum may differ from that of the ensemble mean, $\bar{\mathbf{r}}$ is included as an additional pseudo-absorber with an adjustable coefficient c_{R}, rather than subtracted as a fixed term. The fit matrix becomes

$$\mathbf{B}_{\mathrm{cov}} = [\mathbf{B} \mid \bar{\mathbf{r}}], \tag{17}$$

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

and, taking the weighting matrix to be the covariance of the remaining random error, $\mathrm{Cov}(\boldsymbol{\delta}_{\mathrm{rnd}}) \approx \mathbf{C}$, the fitted parameters follow from the covariance-weighted (GLS) regression

$$\hat{\boldsymbol{\beta}}_{\mathrm{cov}} = \left(\mathbf{B}_{\mathrm{cov}}^{\top}\mathbf{C}^{-1}\mathbf{B}_{\mathrm{cov}}\right)^{-1}\mathbf{B}_{\mathrm{cov}}^{\top}\mathbf{C}^{-1}\boldsymbol{\tau}^{\delta}. \tag{18}$$

The coefficient c_{R} is obtained together with the other fit parameters as the component of $\hat{\boldsymbol{\beta}}_{\mathrm{cov}}$ associated with the $\bar{\mathbf{r}}$ column, and does not require a separate estimation step. Spectral directions that are highly variable in the reference residuals receive less weight, while more stable directions retain a stronger influence on the fit.

2.4 PCA-based residual correction

A complementary strategy is to describe the coherent residual structures explicitly. As a covariance matrix, $\mathbf{C}$ is symmetric with non-negative eigenvalues, and admits the eigendecomposition

$$\mathbf{C} = \sum_{k=1}^{M} \lambda_k \mathbf{v}_k \mathbf{v}_k^{\top}, \qquad \lambda_1 \geq \cdots \geq \lambda_M \geq 0, \qquad \mathbf{v}_i^{\top}\mathbf{v}_j = \delta_{ij}, \tag{19}$$

with orthonormal eigenvectors. The eigenvalue λ_k equals the variance of the residual ensemble along the direction $\mathbf{v}_k$, so the leading K eigenvectors represent the dominant residual patterns. Used as columns of the fit matrix, we refer to them as the PCA components. They are added to the DOAS fit as additional pseudo-absorbers,

$$\mathbf{B}_{\mathrm{PCA}} = [\mathbf{B} \mid \mathbf{v}_1 \cdots \mathbf{v}_K], \tag{20}$$

and the fit is solved by ordinary least squares (uniform weighting), so that PCA-DOAS isolates the effect of fitting the coherent structures out of the residual, independently of any covariance weighting. Coherent structures not represented by the physical cross sections, the polynomial, or the standard pseudo-absorbers are thereby absorbed by the PCA components, which typically improves the sensitivity to weak absorbers (Li et al., 2013; Joiner et al., 2023). Increasing K removes more residual structure but raises the risk of absorbing physical signal; K is therefore chosen as a compromise between residual suppression and retrieval fidelity. For the OClO retrieval we use $K = 4$: this captures the dominant systematic structures carried by the leading eigenvalues of $\mathbf{C}$, while increasing K beyond this value produced no appreciable further reduction of the residual and left the retrieved SCDs essentially unchanged. A known limitation is that the reference ensemble may contain part of the target-gas absorption structure. When this occurs, the PCA components also describe part of the true signal and remove it from the retrieved SCD. We refer to this as target-signal absorption by the PCA components.

2.4.1 Adaptive PCA correction for strong-background absorbers

For trace gases with a strong, spatially extended background such as NO_2, the target-signal absorption by the PCA components can be non-negligible, and the standard correction benefits from an additional adaptive step. A first standard DOAS fit yields an initial target-gas SCD field $\hat{S}_{\mathrm{tgt}}(x,y)$, which is spatially smoothed to retain its large-scale, geophysically coherent component,

$$\tilde{S}_{\mathrm{tgt}}(x,y) = \mathcal{G}_{\boldsymbol{\sigma}} * \hat{S}_{\mathrm{tgt}}(x,y), \tag{21}$$

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

where $\mathcal{G}_{\boldsymbol{\sigma}}$ is a separable Gaussian filter with along- and across-track widths $\boldsymbol{\sigma} = (\sigma_\parallel, \sigma_\perp)$. The smoothed contribution is subtracted from the optical depth,

$$\boldsymbol{\tau}^\dagger(x,y) = \boldsymbol{\tau}^\delta(x,y) - \tilde{S}_{\text{tgt}}(x,y)\,\boldsymbol{\sigma}'_{\text{tgt}}, \tag{22}$$

and the PCA components are computed from the residuals after this subtraction, so that the dominant target-gas structure is removed before the covariance matrix is constructed (Richter, 2026). For the OClO retrieval considered here the target-gas background is weak (e.g. over the equatorial ocean) and the standard PCA correction is used, the adaptive step remains available for strong-background absorbers.

2.5 Covariance–PCA synergy (XDOAS)

XDOAS combines covariance weighting and PCA-based residual correction in a single retrieval. The construction starts from the decomposition of the data error, Eq. (15), into (i) systematic structures of physical or instrumental origin (stray light, residual ISRF distortions, across-track artefacts, undersampling), which appear as deterministic, reproducible patterns across the reference ensemble, and (ii) correlated random noise, which is stochastic but not spectrally white. The two contributions require different treatments: the systematic structures are modelled explicitly as additional pseudo-absorbers, while the correlated random noise is best treated through a covariance-based weighting of the residual.

The systematic part is represented at two complementary levels. The PCA components $\mathbf{v}_k$ are computed from *centred* residuals and therefore describe the dominant fluctuations *about* the mean residual, the mean residual itself is represented separately by $\bar{\mathbf{r}}$. A consistent model of the systematic part is therefore $\mathbf{r}_{\text{sys}} \approx c_{\text{R}}\bar{\mathbf{r}} + \sum_{k=1}^{K} c_k \mathbf{v}_k$, giving the XDOAS fit matrix

$$\mathbf{B}_{\text{X}} = [\mathbf{B}\,|\,\bar{\mathbf{r}}\,|\,\mathbf{v}_1\,\cdots\,\mathbf{v}_K], \tag{23}$$

where $\bar{\mathbf{r}}$ captures the persistent, scene-independent structure that survives the ensemble average (e.g. residual stray-light or ISRF mismatches), while $\{\mathbf{v}_k\}$ capture the scene-dependent patterns that vary in amplitude from pixel to pixel (e.g. across-track ISRF variations or view-angle-dependent stray light).

Once the first K principal directions are included explicitly in the fit matrix, they are no longer part of the remaining random error. The covariance of that remaining error, $\sum_{k=K+1}^{M} \lambda_k \mathbf{v}_k \mathbf{v}_k^\top$, has rank $M-K$ and is therefore singular whenever $K>0$, so its inverse does not exist and a regularised model is required. The variance along the first K directions is replaced by a common background-noise level,

$$\sigma^2_{\text{noise}} = \frac{1}{M-K} \sum_{k=K+1}^{M} \lambda_k, \tag{24}$$

which represents the correlated random noise that remains after the systematic structures have been removed. The regularised covariance and its inverse are

$$\mathbf{C}_{\text{reg}} = \sum_{k=1}^{M} \tilde{\lambda}_k\,\mathbf{v}_k\mathbf{v}_k^\top, \qquad \mathbf{C}_{\text{reg}}^{-1} = \sum_{k=1}^{M} \frac{1}{\tilde{\lambda}_k}\,\mathbf{v}_k\mathbf{v}_k^\top, \qquad \tilde{\lambda}_k = \begin{cases} \sigma^2_{\text{noise}}, & k \le K, \\ \lambda_k, & k > K. \end{cases} \tag{25}$$

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Table 1. Summary of the DOAS retrieval approaches considered in this study. Here, $\mathbf{B}$ is the standard DOAS fit matrix (containing the differential cross sections, the polynomial terms, and the standard pseudo-absorbers of Sect. 2.1), $\mathbf{C}$ is the residual covariance matrix, $\bar{\mathbf{r}}$ is the mean residual of the reference ensemble, $\mathbf{v}_1, \ldots, \mathbf{v}_K$ are the leading PCA components, and $\mathbf{C}_{\mathrm{reg}}$ is the regularised covariance matrix used in XDOAS.

Aspect	Classical DOAS	Covariance-weighted DOAS	PCA-DOAS†	XDOAS
Fit matrix	$\mathbf{B}$	$[\mathbf{B} \,\vert\, \bar{\mathbf{r}}]$	$[\mathbf{B} \,\vert\, \mathbf{v}_1, \ldots, \mathbf{v}_K]$	$[\mathbf{B} \,\vert\, \bar{\mathbf{r}} \,\vert\, \mathbf{v}_1, \ldots, \mathbf{v}_K]$
Residual weighting	Uniform	Covariance-based	Uniform	Regularised covariance-based
Residual treatment	Residuals minimised directly	Correlated residual variability is down-weighted	Coherent residual structures are fitted explicitly	Coherent structures are fitted explicitly remaining variability is covariance-weighted
Main strength	Simple and well established	Reduces the impact of correlated residuals	Represents coherent residual structures	Combines PCA-based correction and covariance weighting on separate parts of the residual
Main limitation	Correlated residuals are not explicitly treated	Residual structures are not explicitly modelled	PCA components may absorb target-gas signal	Requires a representative reference ensemble
Reference	Platt and Stutz (2008)	Theys et al. (2021)	Li et al. (2013)	This study

† PCA-DOAS can be optionally augmented with an adaptive target-signal correction step (Sect. 2.4.1) for retrievals of strong-background absorbers such as NO_2.

The replacement $\lambda_k \to \sigma^2_{\mathrm{noise}}$ for $k \leq K$ is a modelling assumption rather than a consequence of estimation theory, its purpose is to avoid assigning large variances to directions that are already represented explicitly in $\mathbf{B}_{\mathrm{X}}$, while keeping $\mathbf{C}_{\mathrm{reg}}$ positive-definite and invertible. The retrieved parameters follow from the covariance-weighted (GLS) regression

$$\hat{\boldsymbol{\beta}}_{\mathrm{X}} = \left(\mathbf{B}_{\mathrm{X}}^\top \mathbf{C}_{\mathrm{reg}}^{-1} \mathbf{B}_{\mathrm{X}}\right)^{-1} \mathbf{B}_{\mathrm{X}}^\top \mathbf{C}_{\mathrm{reg}}^{-1} \boldsymbol{\tau}^\delta, \tag{26}$$

from which $\hat{\mathbf{S}}$ is read as the first N_g components of $\hat{\boldsymbol{\beta}}_{\mathrm{X}}$. The two mechanisms target distinct parts of the residual and are strictly complementary: the pseudo-absorbers in $\mathbf{B}_{\mathrm{X}}$ remove the systematic structures, while $\mathbf{C}_{\mathrm{reg}}^{-1}$ down-weights the remaining correlated noise.

The connection to the other retrievals follows directly. With $K = 0$ and uniform weighting ($\mathbf{C}_{\mathrm{reg}} = \mathbf{I}_M$), Eq. (26) reduces to the classical DOAS fit of Eq. (6), with $K = 0$ and $\mathbf{C}_{\mathrm{reg}} = \mathbf{C}$, to covariance-weighted DOAS, and with $K > 0$ and uniform weighting, to PCA-DOAS. XDOAS thus encompasses all three as limiting cases, addressing both residual components simultaneously. The degenerate case $K = M$ is avoided in practice, the operating regime is $K \ll M$, where the PCA pseudo-absorbers capture the leading systematic structures while the remaining residual retains sufficient correlated-noise structure for the weighting step to enhance retrieval sensitivity. A full statistical derivation of the four estimators, including the underlying error model and the assumptions (A1)–(A7) on which the empirical covariance rests, is given in Appendix A.

Table 1 summarises the four retrieval approaches and their relationship to XDOAS.

3 OClO retrieval from TROPOMI measurements

The four retrieval approaches summarised in Table 1 are applied to the retrieval of OClO SCDs from TROPOMI/S5P measurements. TROPOMI is a nadir-viewing hyperspectral imaging spectrometer covering the ultraviolet, visible, near-infrared,

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Table 2. DOAS fit parameters used for the OClO retrieval from TROPOMI. The first block lists the absorption cross sections and the Ring spectrum included as pseudo-absorbers in the fit, together with the laboratory temperature of each cross section. The second block lists the global retrieval settings. All cross sections were convolved with the TROPOMI slit function before being used in the fit. The comparison against the S5P-PAL OClO product (Sect. 3.2) uses the same configuration except for the spectral fitting window, which is set to 345.0–389.0 nm.

Absorber	**Reference**	**Temperature**
OClO (target)	Bogumil et al. (2003)	293 K
BrO	Wilmouth et al. (1999)	228 K
NO_2	Vandaele et al. (1998)	220 K
O_4	Thalman and Volkamer (2013)	233 K
O_3	Serdyuchenko et al. (2014)	293 K, 243 K
Ring spectrum	Chance and Kurucz (2010)	—
Global retrieval settings		
Spectral fitting window	344.0–383.0 nm (method evaluation, Sect. 3.1); 345.0–389.0 nm (S5P-PAL comparison, Sect. 3.2)	
Baseline polynomial	Order 5 (6 coefficients)	
Intensity offset	$1/I_0(\lambda)$ pseudo-absorber	
Spike removal threshold	$\kappa = 5.0$ (iterative, against the normalised fit residual)	
Wavelength alignment	Linear shift s and stretch/squeeze q pseudo-absorbers (Eqs. 9–11)	

and shortwave-infrared spectral ranges, with daily global coverage and a nominal ground pixel size of $5.5 \times 3.5\,\mathrm{km}^2$ at nadir (Veefkind et al., 2012). The high signal-to-noise ratio and dense spatial sampling of TROPOMI make it well suited for the retrieval of trace gases across a wide range of atmospheric abundances Theys et al. (2017); Loyola et al. (2018); De Smedt et al. (2018); Garane et al. (2019); Lerot et al. (2021); Liu et al. (2021), and particularly for weak absorbers such as OClO. Because OClO observations are located at high solar zenith angles, where the radiance is low and the residual structures of the DOAS fit are most pronounced, the systematic residual structures discussed in Sect. 2.2 become a leading source of retrieval uncertainty. This makes the OClO retrieval a representative test case for the four residual-treatment approaches compared in this study.

The DOAS fit parameters used in the present study are shown in Table 2. All four retrievals share an identical fit configuration, the same spectral fitting window, baseline polynomial, intensity offset correction, spike-removal threshold, and absorption cross sections, thus any differences in the retrieved SCDs can be directly attributed to the treatment of the fit residual rather than to the underlying fit configuration. The cross sections were convolved with the TROPOMI slit function before being used in the fit, as described in Sect. 2.1.

The retrievals are evaluated on a single representative TROPOMI orbit (orbit 03308, 3 June 2018), and the analysis is divided into two parts. Sect. 3.1 compares the four retrievals with each other using a spectral fitting window of 344.0–383.0 nm,

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

characterising the performance of XDOAS relative to classical DOAS, covariance-weighted DOAS, and PCA-DOAS. Sect. 3.2 then compares the XDOAS retrieval against the S5P-PAL OClO product, using a spectral fitting window of 345.0–$389.0\,\mathrm{nm}$ to match the fitting window used in S5P-PAL product.

3.1 Evaluation of the XDOAS algorithm

Since the four retrievals share identical fit parameters and differ only in how the fit residual is treated, the comparisons in this section isolate the effect of the residual treatment. The four retrievals are evaluated through a set of complementary diagnostics: the retrieval precision in a clean reference sector (Sect. 3.1.1), the distribution of the retrieval uncertainty (Sect. 3.1.2), the across-track behaviour of the retrieved SCDs (Sect. 3.1.3), the global spatial fields (Sect. 3.1.4), and the consistency between the RMS residual and the retrieval uncertainty (Sect. 3.1.5). Each test targets a different aspect of the retrieval, and together they establish the behaviour of XDOAS relative to the three reference retrievals.

3.1.1 Retrieval precision in the equatorial clean sector

The equatorial clean sector (latitudes between $\pm 30°$) is a region where the OClO signal can be assumed to be zero, providing a natural reference in which the retrieval quality can be evaluated without interference from the OClO signal. Within this sector, the standard deviation of the retrieved SCDs quantifies the 1σ scatter that the spectral fit cannot absorb, which is directly representative of the random component of the retrieval uncertainty. Figure 1 shows the running 1σ standard deviation of the retrieved OClO SCD as a function of latitude for the four retrievals.

Classical DOAS shows a background scatter of approximately $2.85 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$, resulting from the combined effect of photon shot noise, residual instrumental structures, and spectral features not captured by the fit, all acting on a uniformly weighted residual. The three advanced retrievals reduce this to a consistent level between 2.0 and $2.5 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$, with a clear minimum near the equator where solar radiation and the signal-to-noise ratio of the measured spectra are highest. The overall reduction relative to classical DOAS is about 20%, depending on latitude. Within the advanced retrievals, covariance-weighted DOAS, PCA-DOAS, and XDOAS show comparable scatter across most of the latitude range, with XDOAS consistently at the lower end.

This behaviour is confirmed by the distribution of the retrieved SCDs over the full clean sector, shown in Fig. 2. Classical DOAS produces a broad distribution with a standard deviation of $2.91 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$ over the full sector, slightly higher than the near-equatorial value of Fig. 1 because the distribution aggregates all latitudes within $\pm 30°$, including those away from the equator where the scatter is slightly larger. The three advanced retrievals yield distributions that are narrower and more sharply peaked, all centred on zero as expected in the absence of OClO absorption. Covariance-weighted DOAS and XDOAS share the same scatter ($2.31 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$) and their distributions are essentially indistinguishable, while PCA-DOAS is slightly broader ($2.37 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$). The absence of any shift of the peak away from zero confirms that none of the advanced retrievals introduces a systematic offset in the retrieved SCDs over the clean sector.

In a region free of OClO absorption, the retrieval scatter is dominated by the random noise of the fit residual. Once the main systematic structures have been removed, either by fitting them as pseudo-absorbers (PCA-DOAS, XDOAS) or by covariance

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

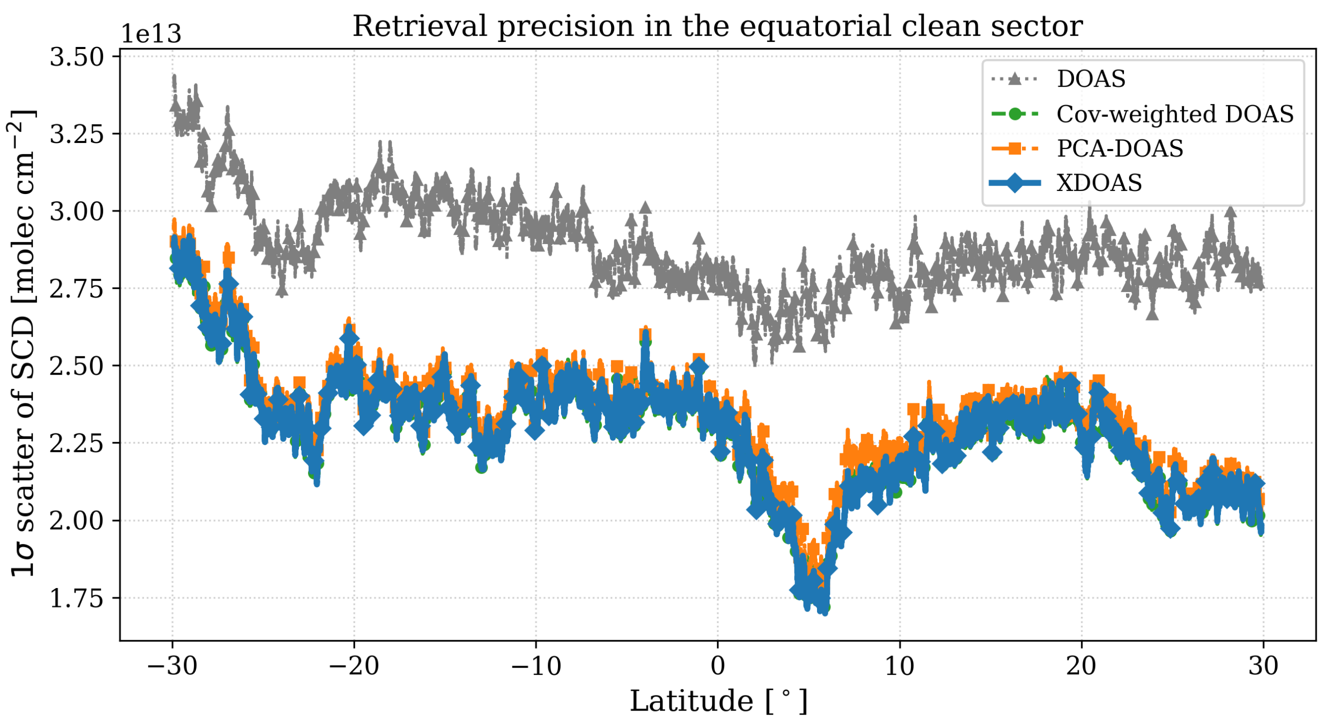

Figure 1. Retrieval precision in the equatorial clean sector ($\pm 30°$). Running 1σ scatter of the retrieved OClO SCD as a function of latitude for the four retrievals.

down-weighting (covariance-weighted DOAS), the remaining scatter is limited by the measurement noise alone, which is the same for all three advanced retrievals. Their agreement on this test is therefore expected and does not imply that the methods are equivalent. The differences between them appear in the tests that probe the quality of the spectral fit more closely, as shown in the following sections.

3.1.2 Retrieval uncertainty

Figure 3 shows the probability density function of the per-pixel retrieval uncertainty σ_{SCD}, derived from the diagonal of the retrieval covariance matrix, evaluated over the equatorial clean sector. The Full Width at Half Maximum (FWHM) of each distribution is used to characterise the typical retrieval uncertainty, as it measures the width around the peak of the distribution without being affected by the asymmetry towards higher uncertainty values.

Classical DOAS shows a broad, asymmetric distribution that peaks around $\sim 3.5 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$ with a FWHM of $9.6 \times 10^{12}\,\mathrm{molec\,cm^{-2}}$. The three advanced retrievals shift this distribution towards lower uncertainty values, with FWHM values of $7.4 \times 10^{12}\,\mathrm{molec\,cm^{-2}}$ (covariance-weighted DOAS), $7.8 \times 10^{12}\,\mathrm{molec\,cm^{-2}}$ (PCA-DOAS), and $7.7 \times 10^{12}\,\mathrm{molec\,cm^{-2}}$ (XDOAS), a reduction of about 20% relative to classical DOAS. The three advanced distributions are nearly overlapping, peaking around $2.2 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$. This reduction is consistent with the improvement in retrieval precision reported in

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

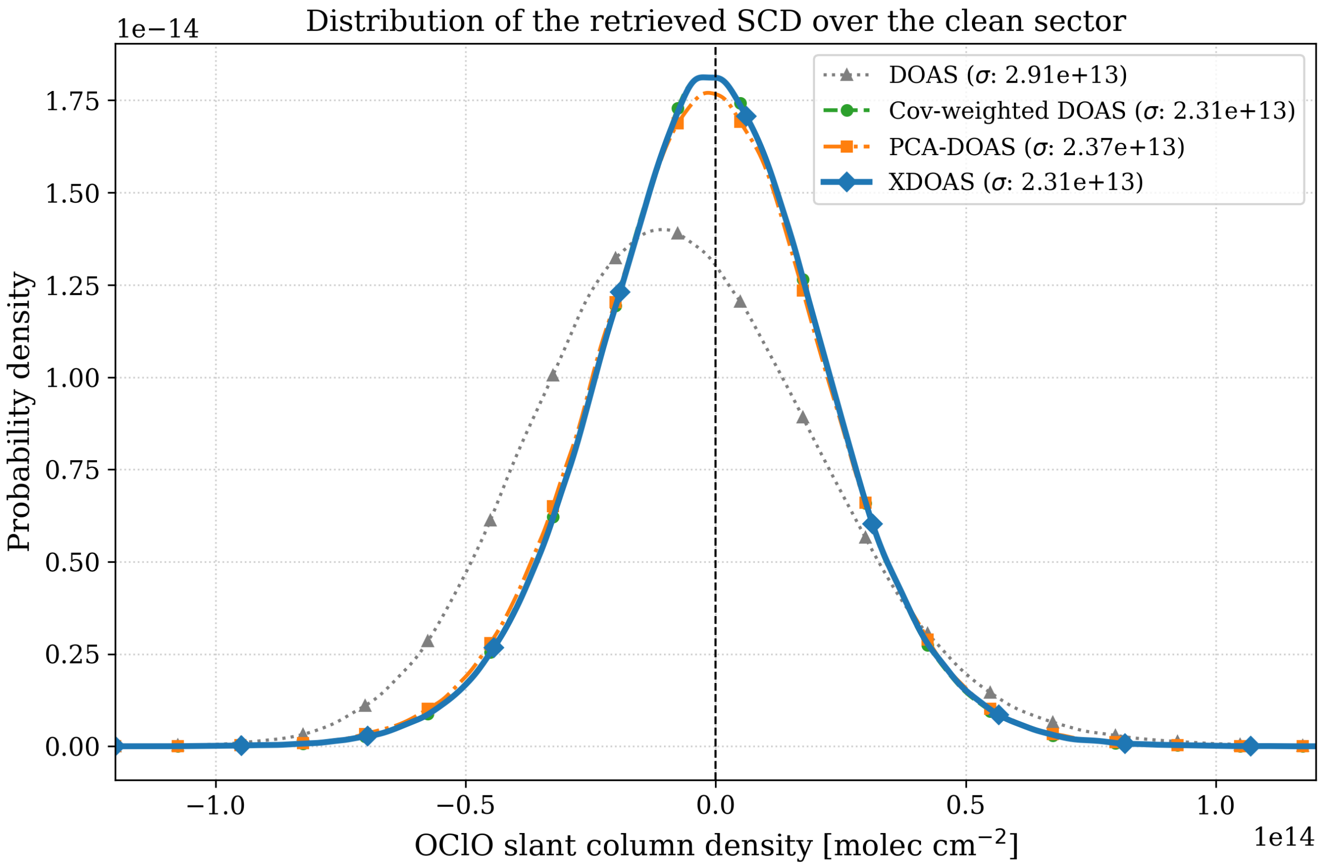

Figure 2. Distribution of the retrieved OClO SCD over the equatorial clean sector ($\pm 30°$) for the four retrievals. The standard deviation of each distribution is given in the legend. The dashed vertical line marks zero SCD, the expected value in the absence of OClO absorption.

Sect. 3.1.1, confirming that the lower retrieval uncertainty corresponds to a real improvement in fit quality rather than an underestimation of the uncertainty.

At first glance, the FWHM values would seem to slightly favour covariance-weighted DOAS, but this result is misleading. The lower FWHM of covariance-weighted DOAS arises from how its retrieval uncertainty is computed. The reported uncertainty of each pixel is the product of two factors: a parameter-dependent term, $\mathrm{diag}[(\mathbf{B}^\top \mathbf{C}^{-1}\mathbf{B})^{-1}]$, which depends on the design matrix and the common covariance but not on the individual residual, and a pixel-dependent variance built from the covariance-weighted residual $\mathbf{r}^\top \mathbf{C}^{-1}\mathbf{r}$ rather than from the unweighted RMS. Because the first factor is nearly identical across the clean-sector pixels, and the second weights the residual by $\mathbf{C}^{-1}$ instead of measuring its unweighted magnitude, the reported uncertainty varies much less from pixel to pixel than the unweighted RMS residual does. The retrieval uncertainty is thus lowered across the whole swath, even for pixels whose unweighted spectral fit quality does not justify it, a point examined in detail in Sect. 3.1.5. The slightly larger FWHM of XDOAS follows directly from the additional pseudo-absorbers

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

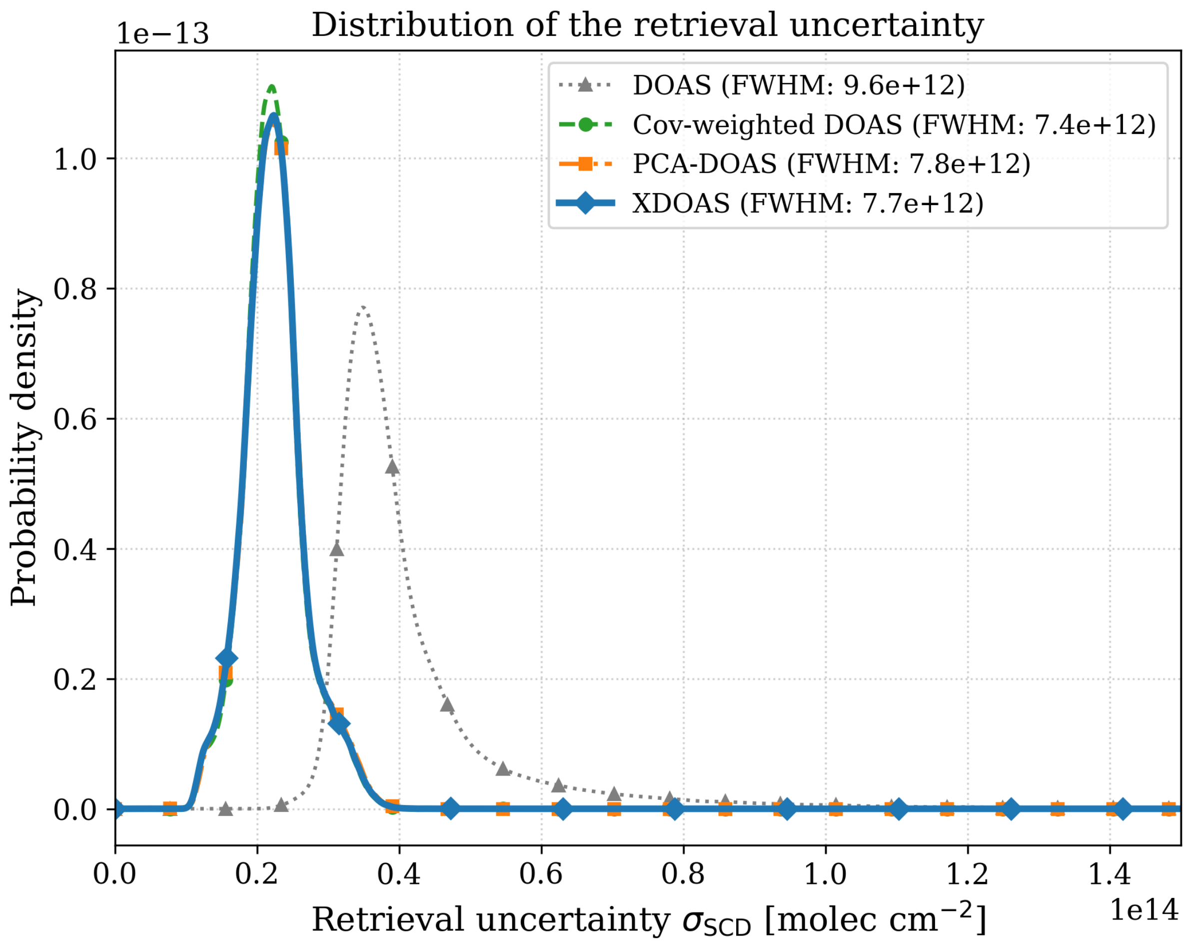

Figure 3. Probability density function of the per-pixel retrieval uncertainty σ_{SCD} for the four retrievals, evaluated over the equatorial clean sector.

introduced by the PCA components, which add degrees of freedom to the fit and broaden the uncertainty distribution by a few
355 percent, while preserving its dependence on the RMS residual of each pixel. PCA-DOAS and XDOAS report nearly identical FWHM because both retrievals remove the same systematic structures through the PCA pseudo-absorbers, and their retrieval uncertainties therefore scale similarly with the RMS residual. The relevant question is therefore not which retrieval produces the lowest FWHM, but which one produces retrieval uncertainties that are consistent with the spectral fit quality of each pixel, a question addressed in Sect. 3.1.5.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

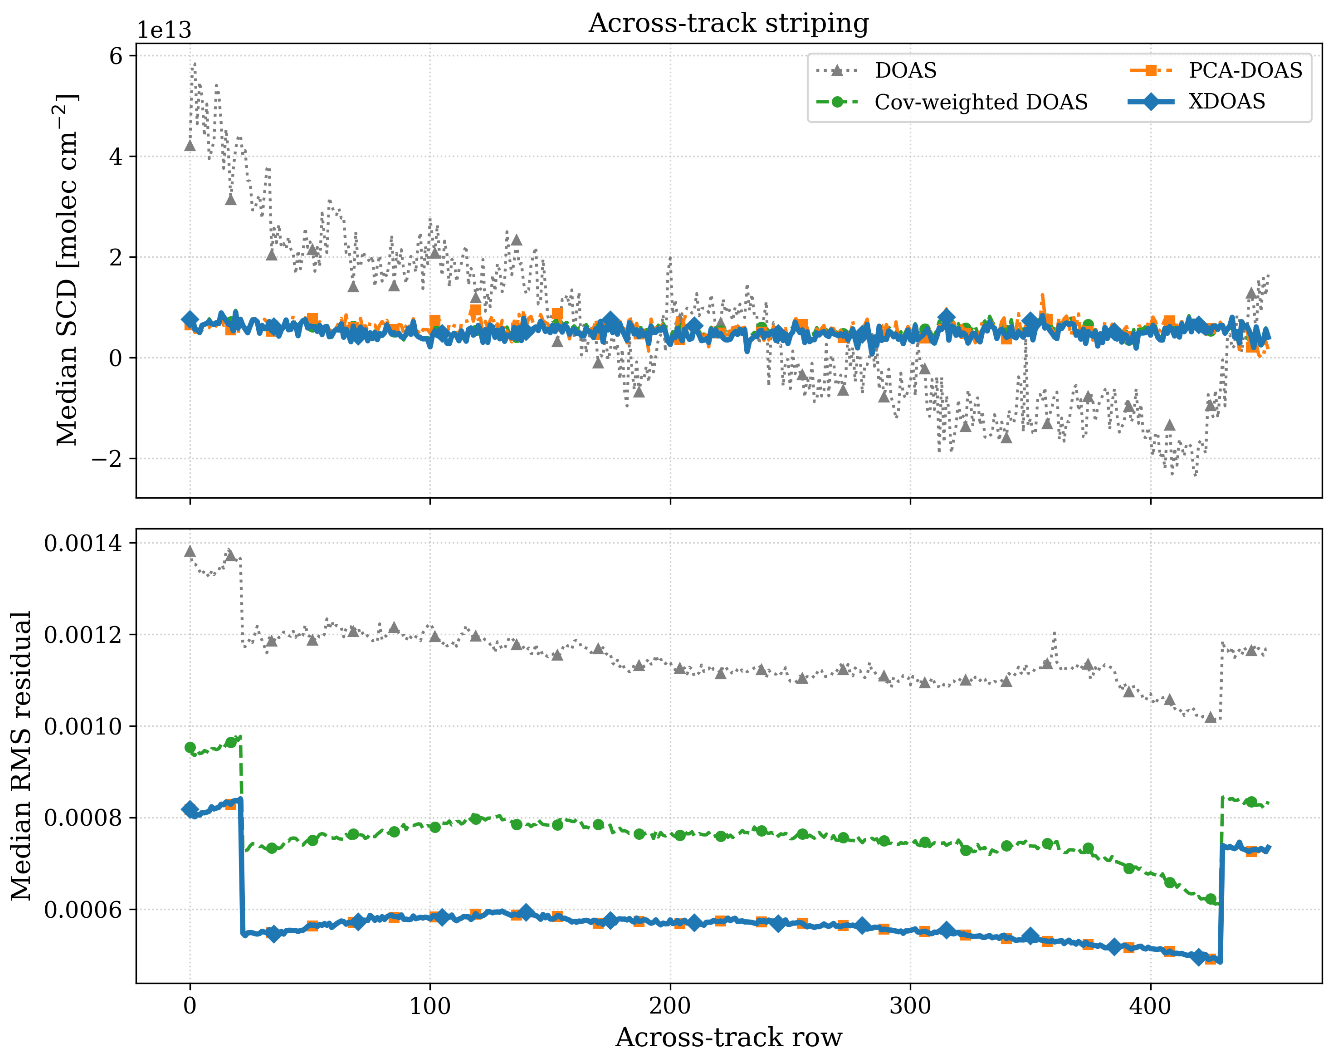

Figure 4. Across-track striping over the full orbit (all along-track scanlines, day side, SZA $\leq 92°$). Upper panel: median retrieved OClO SCD as a function of the across-track row. Lower panel: median RMS residual.

3.1.3 Across-track striping

A well-known challenge in satellite DOAS retrievals from imaging spectrometers is the systematic across-track bias known as striping. It arises because each across-track row of the two-dimensional detector samples the incoming light through a slightly different part of the instrument optics and different detector pixels, and therefore has its own instrument spectral response function, radiometric response, and stray-light contribution. Figure 4 shows the median retrieved OClO SCD (upper panel) and the median RMS residual (lower panel) as a function of the across-track row for the four retrievals.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Classical DOAS shows a pronounced striping pattern in the retrieved SCDs, with median values reaching $\sim 5\times10^{13}\,\mathrm{molec\,cm^{-2}}$ at the swath edges and a systematic across-track gradient that drifts towards negative values over the second half of the swath. This is the characteristic signature of uncorrected detector-row and ISRF variations, which bias the fit differently for each across-track position. The three advanced retrievals flatten this pattern to a nearly constant level of $\sim 0.5\times10^{13}\,\mathrm{molec\,cm^{-2}}$, within $\pm1\times10^{13}\,\mathrm{molec\,cm^{-2}}$, across the whole swath, removing the striping almost entirely. In the SCD field, covariance-weighted DOAS, PCA-DOAS, and XDOAS behave almost identically, with no residual striping visible above the pixel-to-pixel scatter.

The differences among the advanced retrievals emerge in the RMS residual (lower panel), which separates them clearly into three levels. Classical DOAS lies highest, at about 1.1–1.4×10^{-3}. Covariance-weighted DOAS occupies an intermediate level around 7–9×10^{-4}. PCA-DOAS and XDOAS reach the lowest RMS residual, around 5–6×10^{-4}, and are almost indistinguishable across the swath. These differences are explained by the residual-treatment mechanisms described in Sect. 2. Covariance-weighted DOAS reduces the influence of the spectral directions that dominate the residual covariance, but does not remove them from the residual, whereas PCA-DOAS and XDOAS include the dominant cross-track patterns as pseudo-absorbers and fit them out of the residual entirely. The lower RMS residual of the PCA-based retrievals is the direct consequence of fitting these structures explicitly rather than leaving them in the residual. The steps in the RMS residual at the extreme swath edges, common to all retrievals, correspond to the outermost detector rows, where the spectral smile effect is largest and the ISRF characterisation is least accurate.

3.1.4 Spatial behaviour

Figure 5 shows the global distribution of the retrieved OClO SCD (top row), the RMS residual (middle row), and the retrieval uncertainty (bottom row) for the four retrievals.

Classical DOAS (leftmost column) shows two distinct artefacts in the SCD field. The first is across-track striping at high northern latitudes, caused by the row-to-row differences in the detector response discussed in Sect. 3.1.3. The second is a large-scale gradient in the background SCD along the orbit, which changes sign between the two halves of the orbit: the retrieved background is systematically positive over one half and negative over the other. This pattern follows the change in viewing geometry and solar illumination along the orbit and is therefore instrumental in origin, unrelated to any real OClO signal. Covariance-weighted DOAS (second column) reduces both artefacts and gives a visibly cleaner SCD field, but its RMS-residual map is still irregular: the reweighting lowers the variability without removing the systematic structures, which stay in the residual. PCA-DOAS (third column) reduces them further by fitting the main systematic structures out of the residual, giving cleaner and smoother RMS-residual and retrieval-uncertainty maps. XDOAS (rightmost column) gives an SCD field that is essentially identical to PCA-DOAS by eye, while its RMS-residual and retrieval-uncertainty maps are the smoothest of the four. In the SCD field alone, PCA-DOAS and XDOAS cannot be distinguished, the difference between them appears only in the RMS–uncertainty consistency examined in Sect. 3.1.5.

The point of this comparison is that the geophysical signal is preserved while the instrumental artefacts are removed. The main features of the OClO field, the strong activated-chlorine signal over the southern polar vortex and the weak background at

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

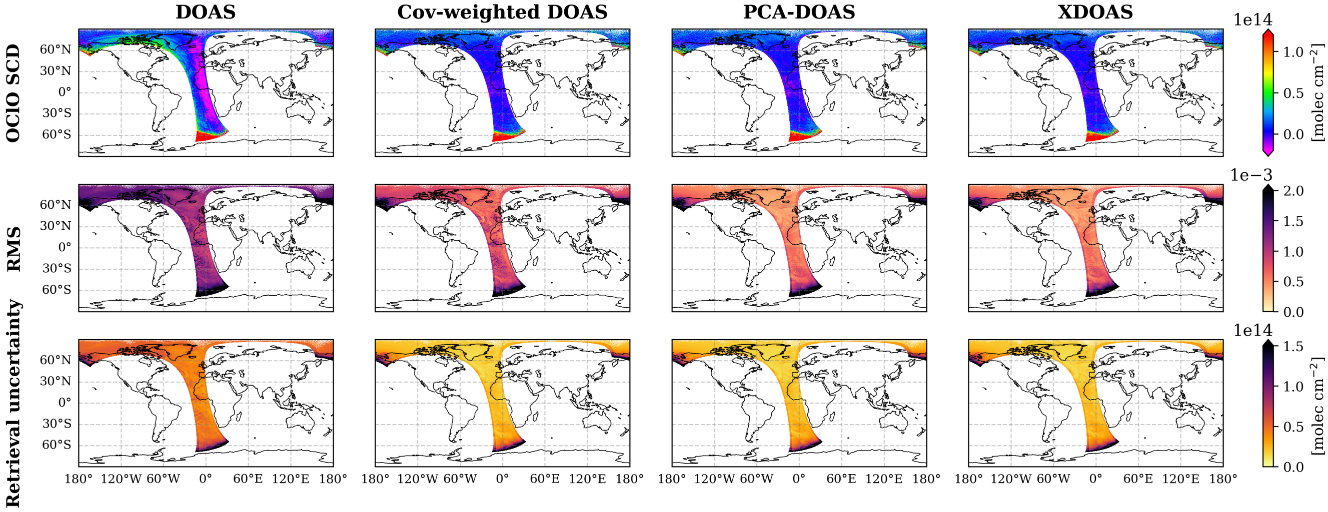

Figure 5. Global spatial behaviour of the four retrievals. Rows from top to bottom: retrieved OClO SCD, RMS residual, and per-pixel retrieval uncertainty. Columns from left to right: classical DOAS, covariance-weighted DOAS, PCA-DOAS, and XDOAS.

mid-latitudes, are kept by all four retrievals. Only the PCA-based retrievals, however, combine clean suppression of the striping and the along-orbit gradient with smooth RMS-residual and uncertainty maps.

3.1.5 Consistency between RMS residual and retrieval uncertainty

The tests so far have shown that the three advanced retrievals lower the measurement noise (Sect. 3.1.1), narrow the retrieval-uncertainty distribution (Sect. 3.1.2), remove the across-track striping (Sect. 3.1.3), and give similarly clean spatial fields (Sect. 3.1.4). On all of these tests, PCA-DOAS and XDOAS are almost identical. The test presented here is the one that separates them, and it is therefore the most important comparison of the internal evaluation. It verifies whether the retrieval uncertainty reported by each method remains consistent with the spectral fit residual that is actually observed. As in the precision and uncertainty tests, it is evaluated over the equatorial clean sector, where the absence of OClO signal means that both the RMS residual and the retrieval uncertainty are determined by the noise of the fit alone.

Unlike the spectral-fit R^2 of a standard DOAS fit, where a value closer to unity indicates a better fit, the R^2 of this consistency test must be interpreted differently. Here the two axes are not the model and the data, but the observed residual and the uncertainty that each method reports for it. A value driven towards unity does not indicate a better retrieval; it indicates that the reported uncertainty has lost its independence from the residual and has become a deterministic rescaling of it. The quantity of interest is therefore the balance between the two, not the magnitude of R^2. This R^2 is specific to the consistency test and does not qualify the RMS-residual comparisons of Sects. 3.1.3–3.1.4, where a lower residual does indicate a genuinely better spectral fit: the RMS residual measures how well the model reproduces the spectrum, whereas the R^2 here measures only how strongly the reported uncertainty is tied to that residual.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

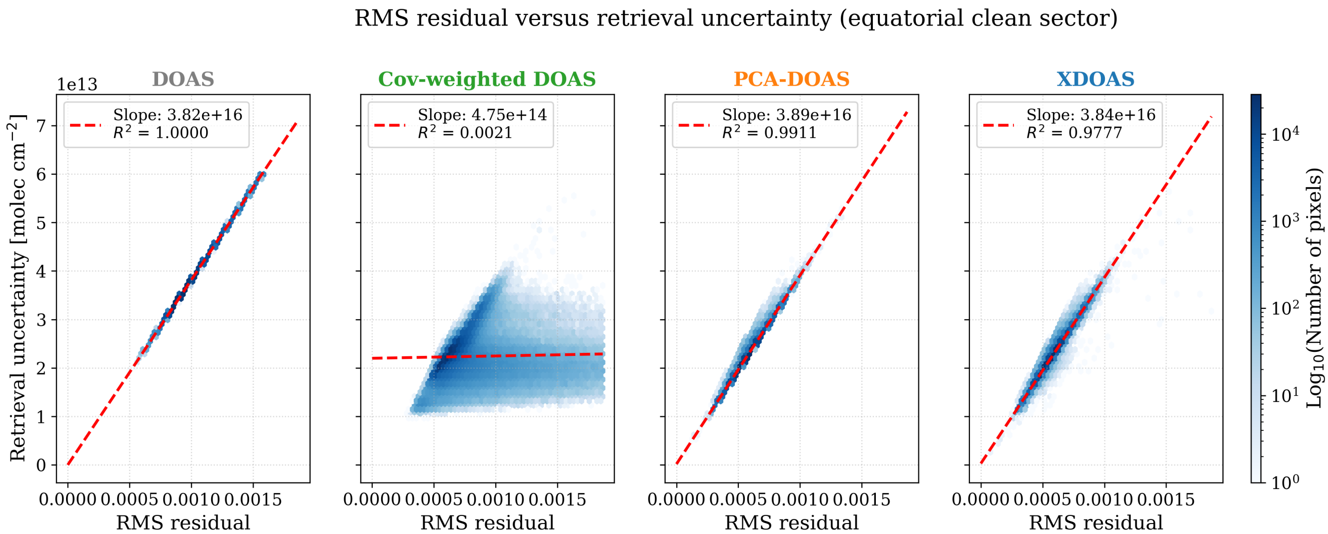

Figure 6. Density of pixels in RMS residual (unweighted) versus retrieval uncertainty over the equatorial clean sector, for the four retrievals. The red dashed line shows the linear regression, with the slope and R^2 in the legend.

Two quantities are compared. The RMS residual on the x–axis is the unweighted root-mean-square of the spectral fit residual, $\mathrm{RMS} = \sqrt{\sum_i r_i^2/N}$, computed in the same way for all four retrievals. It measures directly how much of the measured optical depth the fit failed to explain. The retrieval uncertainty on the y–axis is the square root of the diagonal of the fit covariance matrix. For classical DOAS and PCA-DOAS this uncertainty is computed from the same unweighted residual, whereas for covariance-weighted DOAS and XDOAS it is computed from the residual weighted by the inverse noise covariance $\mathbf{C}^{-1}$. For the covariance-based retrievals, the two axes therefore measure the residual in different ways. The dispersion around the regression line then measures how closely the retrieval uncertainty of each pixel follows its unweighted RMS residual: a small but non-degenerate spread indicates an uncertainty that stays consistent with the residual actually observed. This spread arises because the reported per-pixel uncertainty depends not only on the residual magnitude but also on the geometric sensitivity term $\mathrm{diag}[(\mathbf{B}^\top\mathbf{C}^{-1}\mathbf{B})^{-1}]$, which varies from pixel to pixel; two pixels with the same RMS residual can therefore carry different, correctly sized uncertainties. There is accordingly no target value of R^2, and a lower R^2 is not by itself better; the criterion for a correctly sized uncertainty is $\chi^2_{\mathrm{red}} \approx 1$ (Sect. 3.2.5), not the magnitude of R^2. Figure 6 shows the density of pixels for the four retrievals, together with the slope and the R^2 of the regression.

The most important quantity in this test is not the slope of the regression, which is essentially the average geometric sensitivity of the fit, nor the presence of a linear correlation, but the dispersion of the pixels around the regression line, summarised by R^2.

Classical DOAS (Fig. 6, leftmost panel) shows an exact linear correlation, with a slope of 3.82×10^{16} and $R^2 = 1.0000$. This perfect correlation is not a sign of quality but a trivial consequence of the calculation. In the absence of covariance weighting, the retrieval uncertainty is computed from the same unweighted residual that defines the RMS, and the two axes therefore carry

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

the same information. Classical DOAS is the baseline case, in which the reported uncertainty follows the RMS residual by definition.

Covariance-weighted DOAS (second panel) loses this consistency completely. The points scatter away from the regression line and collapse into a nearly flat band, with R^2 dropping to 0.002. The regression slope falls to 4.75×10^{14}, almost two orders of magnitude below that of the other retrievals, because the reported uncertainty has become nearly constant across pixels and no longer depends on the RMS residual at all. In the clean sector, where there is no OClO signal, the covariance weighting lowers the uncertainty of every pixel to a similar value regardless of its actual residual, and a pixel with a large RMS residual can be assigned the same uncertainty as a pixel with a small one. This is the origin of the smaller FWHM found for covariance-weighted DOAS in Sect. 3.1.2: the lower FWHM is not a sign of better precision, but the result of a retrieval uncertainty compressed to a nearly constant value that no longer tracks the spectral fit quality of each pixel. This is the main reason not to use covariance-weighted DOAS on its own.

PCA-DOAS (third panel) recovers a linear correlation, with a slope of 3.89×10^{16} and $R^2 = 0.991$. In this consistency test, however, a near-perfect correlation is not the desired outcome. The retrieval uncertainty should not be a deterministic rescaling of the RMS residual; it should also reflect how well each fit constrains the slant column, and therefore retain some pixel-to-pixel variability that is independent of the residual magnitude alone. Because PCA-DOAS adds the leading residual components as pseudo-absorbers in an unweighted fit, the same components that lower the residual also lower the reported uncertainty together with it. The correlation tightens towards unity precisely because the uncertainty has lost its independence from the residual, which is the signature of the PCA pseudo-absorbers having absorbed part of the random noise, not only the systematic structures. On average the fit is correct, but the retrieval uncertainty is no longer an independent measure of the fit quality.

XDOAS (rightmost panel) shows the behaviour expected of a consistent fit, a clean linear correlation with a comparable slope (3.84×10^{16}, close to PCA-DOAS because both retrievals use the PCA pseudo-absorbers to fit the systematic structures out of the residual), but with $R^2 = 0.977$, slightly below that of PCA-DOAS. This lower value is not a worse one: it reflects that the random noise is no longer absorbed into the fitted components, leaving a moderate pixel-to-pixel spread. This test alone, however, cannot establish whether that spread corresponds to a correctly sized uncertainty or is merely additional scatter; that question is settled by the independent reduced-χ^2 calibration of Sect. 3.2.5 (Fig. 12), which compares the reported uncertainty against the observed clean-sector scatter and yields $\chi^2_{\mathrm{red}} \approx 1$ for XDOAS. The regularisation of the covariance along the directions of the PCA pseudo-absorbers (Eq. 25) keeps the weighted uncertainty close to the unweighted RMS residual, preventing the decoupling seen in covariance-weighted DOAS, while the explicit fitting of the systematic structures prevents the rigid, over-tight correlation of PCA-DOAS. The retrieval uncertainty of XDOAS is neither compressed to a constant value nor fixed to the RMS residual, but follows the spectral fit quality of each pixel with the dispersion that a real fit should have.

This balanced behaviour is the main characteristic of XDOAS and the reason to prefer it over either of the individual approaches. The key point is that XDOAS achieves this consistency without sacrificing performance on any other test: it matches the best individual retrieval everywhere else, and is the only method that also keeps the retrieval uncertainty consistent with the observed scatter, as verified quantitatively by the reduced-χ^2 calibration of Sect. 3.2.5. Its precision is comparable to

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

that of covariance-weighted DOAS and PCA-DOAS (Sect. 3.1.1), its retrieval-uncertainty narrowing is comparable to that of PCA-DOAS (Sect. 3.1.2), and it removes the striping as effectively as PCA-DOAS does (Sect. 3.1.3). On this test, which measures whether the retrieval uncertainty stays consistent with the observed residual, the three advanced retrievals separate clearly. Covariance-weighted DOAS gives uncertainties that are independent of the residual ($R^2 = 0.002$), PCA-DOAS gives uncertainties that are tied too rigidly to the residual because the PCA pseudo-absorbers fit out part of the random noise ($R^2 = 0.991$), and XDOAS gives uncertainties that follow the spectral fit quality of each pixel with a moderate pixel-to-pixel spread ($R^2 = 0.977$). Read against what a consistent uncertainty estimate requires, rather than as a goodness-of-fit ranking, these values show that XDOAS provides the most balanced relationship between residual magnitude and retrieval uncertainty. Already at this stage, before any comparison with the S5P-PAL product, this balance points to the quality of XDOAS.

3.2 Comparison against the S5P-PAL product

The evaluation of Sect. 3.1 has shown that XDOAS is the best behaved of the four retrievals on its own terms: it combines the ability of PCA-DOAS to fit the systematic structures out of the residual with the lower measurement noise common to all advanced retrievals, while avoiding the two problems seen in the RMS–uncertainty relationship: the underestimated uncertainties of covariance-weighted DOAS and the overfitting of PCA-DOAS. These advantages are now tested against an independently developed reference, by comparing the XDOAS retrieval with the S5P-PAL OClO product (Meier et al., 2024).

For the comparison to be consistent, both products use the same DOAS fit parameters. The XDOAS retrieval is run on the same TROPOMI input data as the S5P-PAL product, with the same spectral fitting window (345.0–$389.0\,\mathrm{nm}$), the same baseline polynomial (order 5), and the same cross sections. Both products are filtered in the same way, by solar zenith angle (SZA $\leq 92°$) and valid values only, with no quality-flag screening applied to either, so that the two algorithms are compared on equal terms.

The two retrievals differ in how the systematic residual structures are treated. The S5P-PAL product treats them after the fit: it includes empirical pseudo-absorbers derived from the mean residual to correct the negative biases at high latitudes, together with an explicit destriping and offset correction over the equatorial Pacific (Meier et al., 2024). XDOAS does not apply any of these post-processing steps. The same systematic structures that these corrections address, including the across-track striping, are fitted out of the residual within the DOAS fit itself, through the PCA pseudo-absorbers and the regularised covariance. The comparison therefore contrasts two ways of treating the systematic residual structures: an empirical post-correction in S5P-PAL, and an in-fit treatment in XDOAS. Any systematic difference between the two products can then be attributed to this difference in the residual treatment and not to the fit parameters.

3.2.1 Global agreement

Figure 7 shows the pixel-to-pixel correlation between the XDOAS and S5P-PAL OClO SCDs, for the full orbit and separated into two regions: the equatorial clean sector, where there is no OClO signal, and the southern polar vortex, where the signal is strong.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

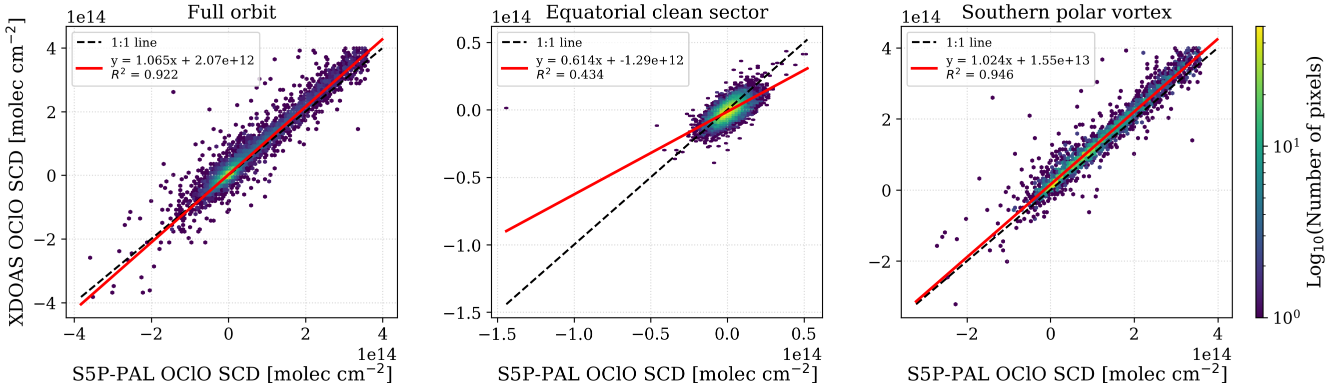

Figure 7. Pixel-to-pixel correlation between the XDOAS and S5P-PAL OClO slant column densities. Left: full orbit. Centre: equatorial clean sector ($|\text{lat}| \leq 30°$). Right: southern polar vortex ($\text{lat} \leq -55°$). The dashed line is the 1:1 reference, the red line is the linear regression.

Over the full orbit, the two retrievals agree with an R^2 of 0.922, a regression slope of 1.065, and an intercept of $2.07 \times 10^{12}\,\text{molec}\,\text{cm}^{-2}$. This global slope above the 1:1 line does not, however, describe the agreement in the same way across the orbit, as the two regional correlations show. In the southern polar vortex, where the OClO signal is strong, the agreement is almost exact, the slope is 1.024 and the R^2 is 0.946, so XDOAS and S5P-PAL recover the same SCDs where the signal is largest, which is the most relevant condition for OClO. In the equatorial clean sector, where there is no OClO signal, the slope flattens to 0.61 and the R^2 drops to 0.434. This is expected in the absence of signal, the correlation is between the scatter of one product and the scatter of the other around zero, not between a real signal common to both, so the slope and the R^2 lose their usual meaning. The empirical offset and destriping that S5P-PAL applies to the background, which XDOAS does not, add to this difference in the clean sector. The global slope of 1.065 is therefore not a uniform 6.5% difference in the retrieved SCDs, but the result of combining the signal-free background, where the correlation decreases, with the vortex, where the two products agree almost exactly.

The distribution of the per-pixel difference $\Delta\text{SCD} = \text{XDOAS} - \text{S5P-PAL}$, shown in Fig. 8, confirms that the absolute difference between the two products is small. The distribution peaks close to zero, with a median difference of $1.05 \times 10^{12}\,\text{molec}\,\text{cm}^{-2}$, a mean of $2.15 \times 10^{12}\,\text{molec}\,\text{cm}^{-2}$, and a standard deviation of $7.38 \times 10^{12}\,\text{molec}\,\text{cm}^{-2}$. The median difference is therefore about one order of magnitude below the typical retrieval uncertainty of each product ($\sim 2 \times 10^{13}\,\text{molec}\,\text{cm}^{-2}$, Sect. 3.1.2) and about two orders of magnitude below the range of the polar OClO signal ($\sim 10^{14}\,\text{molec}\,\text{cm}^{-2}$). For reference, a Gaussian of equal mean and standard deviation is also shown, the measured distribution does not need to agree with it, and indeed it is slightly asymmetric towards positive values, with the mean larger than the median, which is the signature of the small excess of positive values. This small offset comes from the slightly larger fraction of systematic residual structure that XDOAS fits out, and it is well below the noise of each retrieval.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

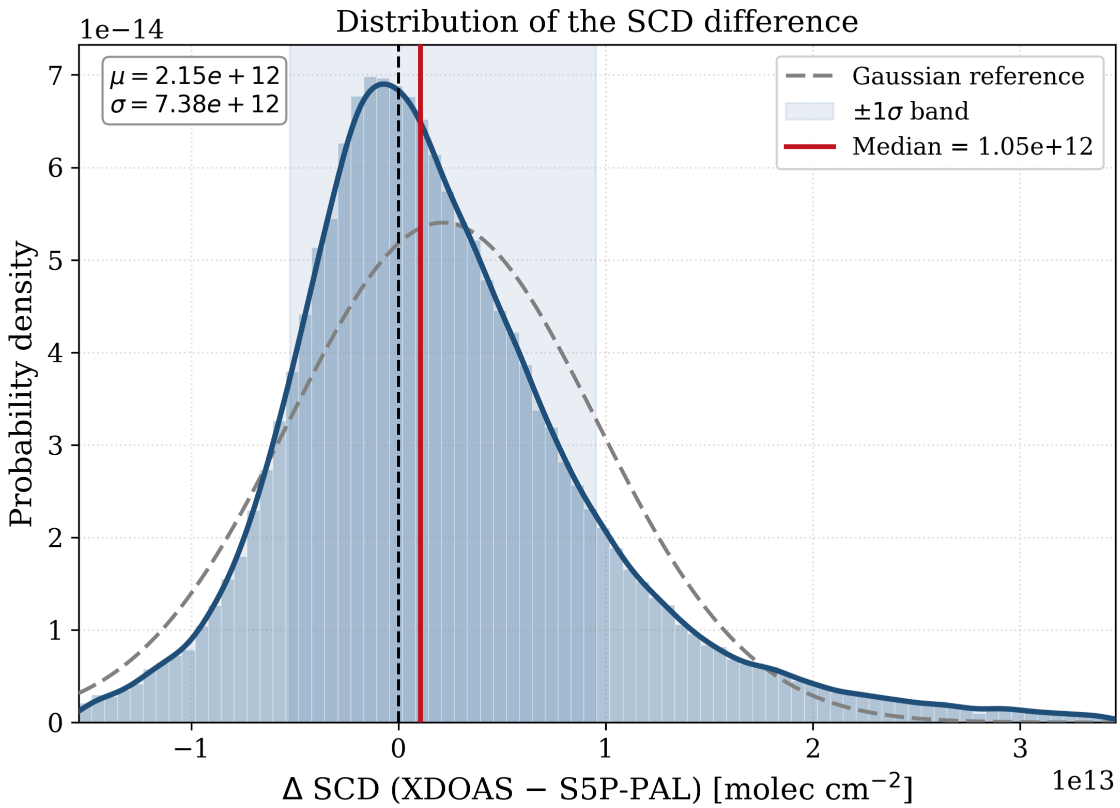

Figure 8. Distribution of the per-pixel difference between the XDOAS and S5P-PAL OClO slant column densities. The dashed line is a Gaussian of equal mean and standard deviation, shown for reference.

3.2.2 Latitudinal consistency

The global agreement does not show where, along the latitude structure of the OClO field, the small differences arise. Figure 9 shows the zonal-mean OClO SCD as a function of latitude for both retrievals, with shaded bands for the 1σ scatter of the pixels in each latitude bin.

The two zonal-mean curves overlap across the whole latitude range. In the mid-latitude background (about $-50°$ to $+50°$), both retrievals scatter around zero with a 1σ band of a few $10^{12}\,\mathrm{molec\,cm^{-2}}$ that is the same for the two products, which confirms that the small positive offset seen in the global agreement (Sect. 3.2.1) is not located in any particular latitude band, but is the uniform result of the difference in the residual treatment. Over the Antarctic polar vortex (south of about $-55°$), the two curves increase together to values above $2 \times 10^{14}\,\mathrm{molec\,cm^{-2}}$, with XDOAS and S5P-PAL following each other within their 1σ bands across the full SCD range. This shows that the geophysical content of the S5P-PAL product is preserved by XDOAS where the OClO signal is largest.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

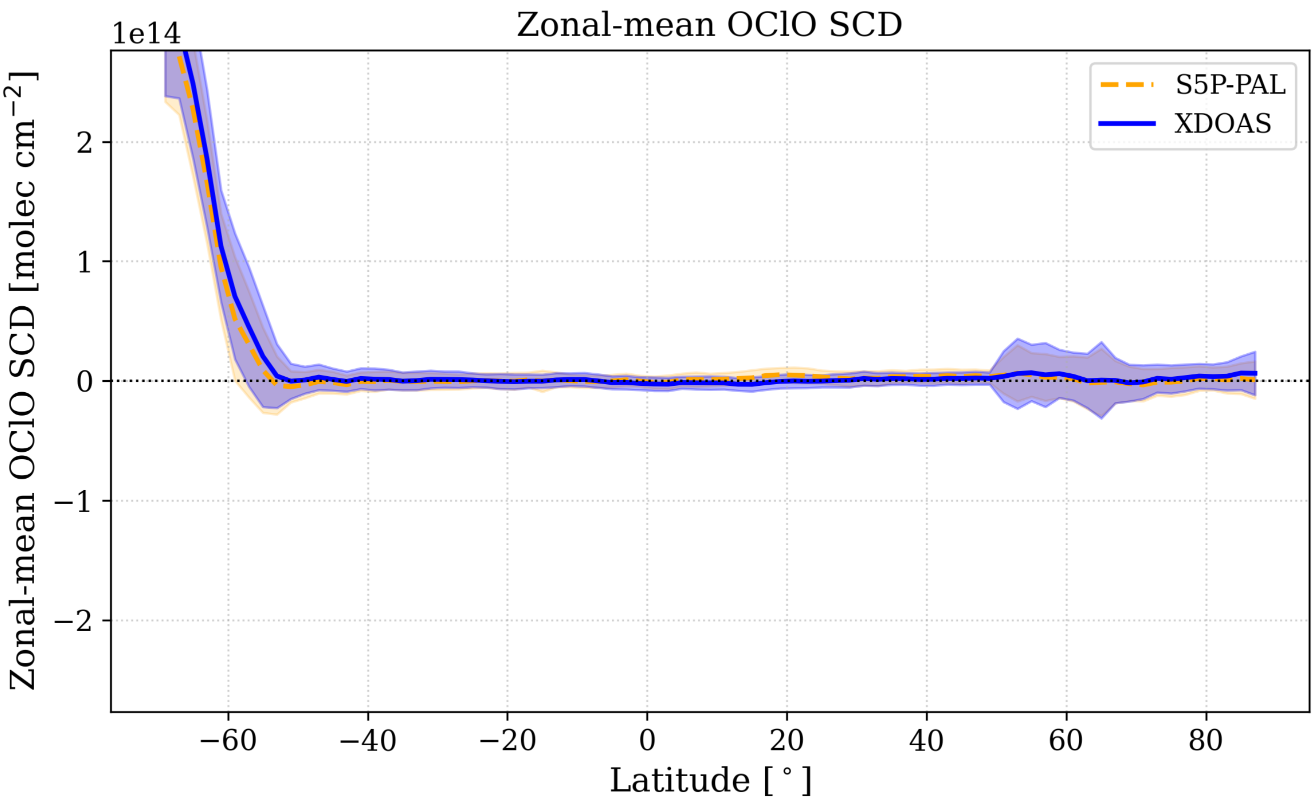

Figure 9. Zonal-mean OClO slant column density as a function of latitude for XDOAS (blue) and S5P-PAL (orange). The shaded bands show the 1σ scatter of the pixels in each latitude bin.

535 A further point can be made in the northern high latitudes (about $+50°$ to $+85°$), where the orbit crosses the day-night boundary and the radiance drops sharply. Here the XDOAS 1σ uncertainty is visibly larger than that of the S5P-PAL product, which at first sight might suggest that XDOAS performs worse. The opposite is true. The larger XDOAS 1σ uncertainty correctly propagates the larger residual scatter near the day-night boundary, in line with the RMS–uncertainty consistency of Sect. 3.1.5. The smaller S5P-PAL 1σ uncertainty is consistent with the pattern seen in the uncertainty maps of Fig. 11,
540 where the S5P-PAL uncertainty appears smaller than the observed scatter of these pixels. The two retrievals agree within their uncertainties, but the two are not equivalent: the XDOAS 1σ uncertainty tracks the observed scatter of the slant columns in this challenging illumination regime, whereas the S5P-PAL uncertainty appears smaller than the observed scatter would suggest.

3.2.3 Stability under extreme optical paths

The solar zenith angle is a direct measure of the geometric difficulty of the retrieval: at high SZA the light path through the
545 atmosphere is long, the signal-to-noise ratio of the measured radiance drops, and any systematic structure not described by the

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

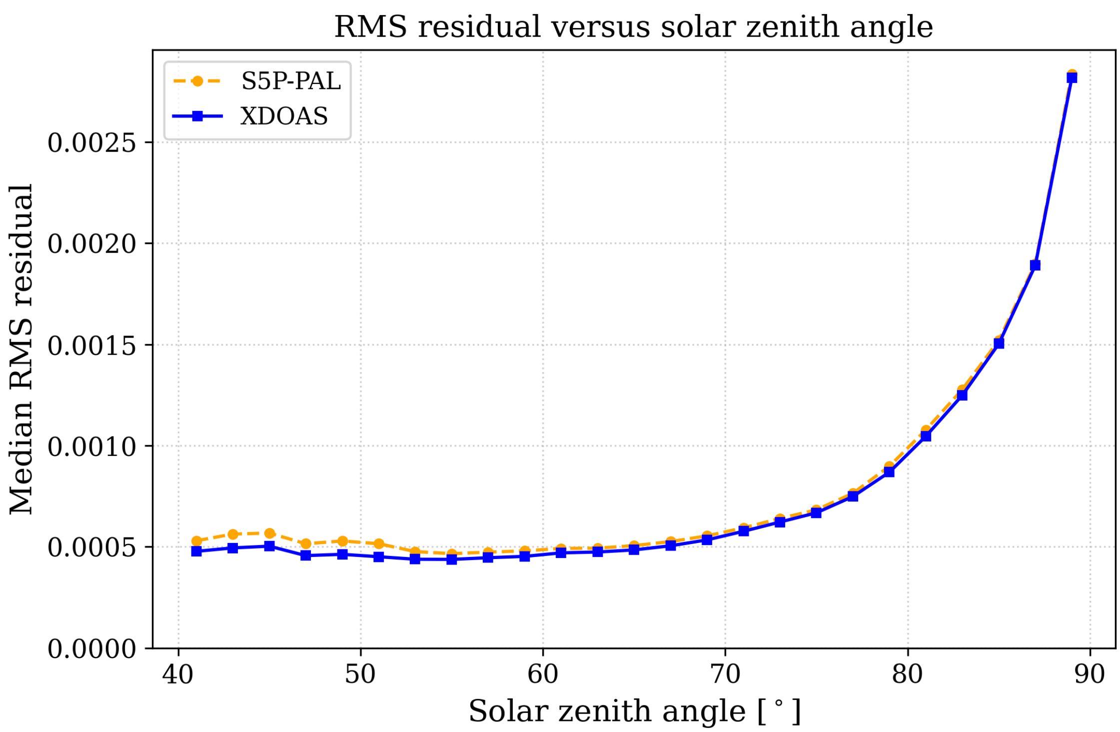

Figure 10. Median RMS residual as a function of solar zenith angle for XDOAS (blue) and S5P-PAL (orange).

fit model contributes more strongly to the residual. Figure 10 shows the median RMS residual as a function of SZA for both XDOAS and S5P-PAL.

The two curves are close up to about $\mathrm{SZA} = 75°$, where the RMS residual is dominated by photon shot noise and both retrievals are stable. Beyond about $80°$, where the radiance drops by more than one order of magnitude and the residual is increasingly dominated by stray-light and ISRF structures, the XDOAS curve stays below the S5P-PAL curve, with the difference growing as the SZA approaches the day-night boundary. By fitting the systematic residual structures out through the PCA pseudo-absorbers and down-weighting the random part through the regularised covariance, XDOAS maintains the RMS residual lower exactly where the S5P-PAL product is most affected by systematic structures. Since polar chlorine activation is observed mostly under low-illumination conditions, the lower RMS residual of XDOAS in this range is directly relevant for OClO measurements.

3.2.4 Spatial distribution and differences

Figure 11 shows the spatial maps of the OClO SCD (top row), the RMS residual (middle row), and the retrieval uncertainty (bottom row) for XDOAS (left column) and S5P-PAL (middle column), together with their difference (right column).

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

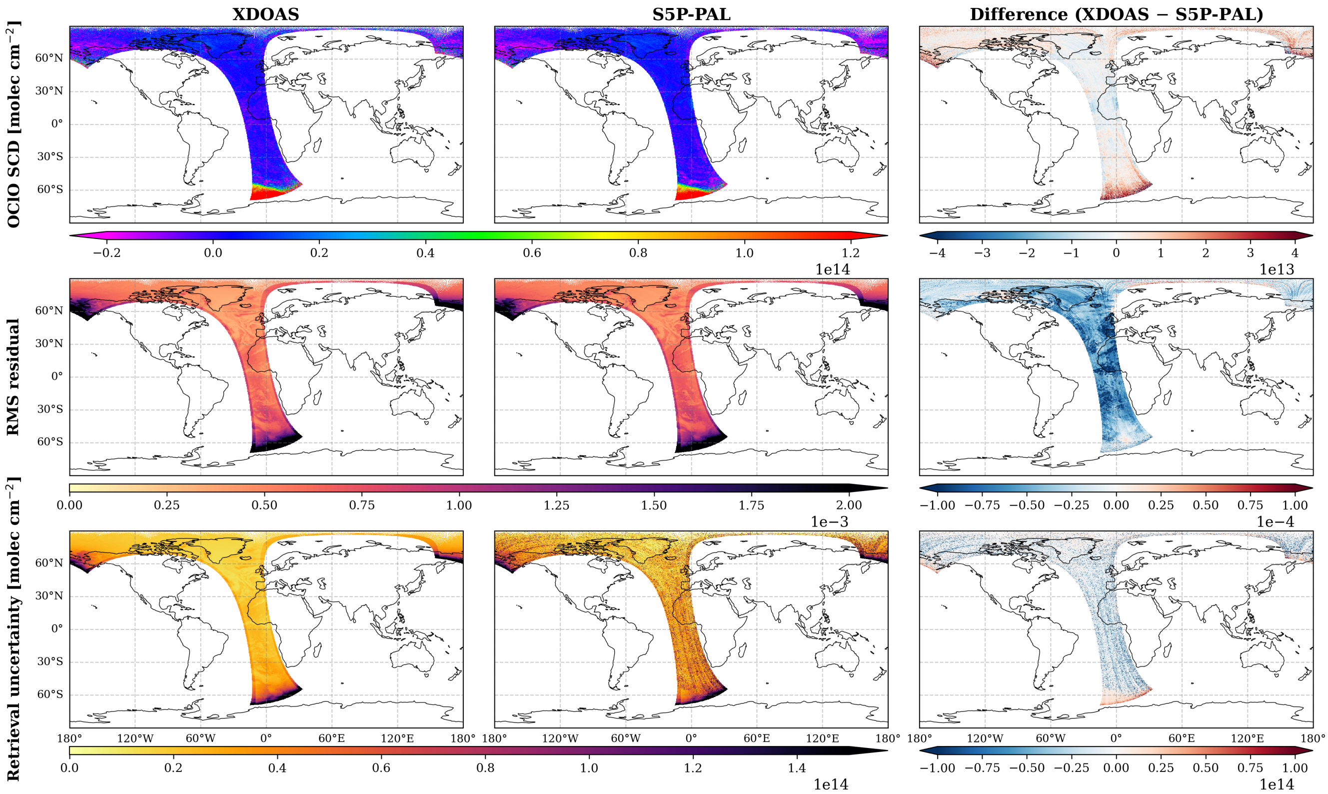

Figure 11. Spatial distribution and differences between XDOAS and S5P-PAL. Columns from left to right: XDOAS, S5P-PAL, and the difference (XDOAS−S5P-PAL). Rows from top to bottom: OClO SCD, RMS residual, and retrieval uncertainty.

The SCD fields of XDOAS and S5P-PAL show no visible difference, both products reproduce the strong OClO signal of the activated Antarctic vortex and the low background at lower latitudes. The difference map (top-right) shows a small positive signal over the polar vortex, consistent with the slope above the 1:1 line in Fig. 7, and a near-zero background elsewhere, with no systematic latitudinal or longitudinal structure outside the polar regions. This is the expected behaviour for two retrievals that share the same fit parameters and differ only in the residual treatment: the geophysical content is recovered in the same way, and the residual treatment changes the result only where the systematic structures contribute significantly.

The RMS residual maps (middle row) show a systematic improvement of XDOAS over S5P-PAL. The XDOAS RMS field is lower than the S5P-PAL field across the whole orbit, and the difference is most visible at high northern latitudes where the orbit crosses the day-night boundary. The difference map (middle-right) is mostly negative (XDOAS below S5P-PAL), which confirms that XDOAS fits out a larger fraction of the systematic residual structure. This is consistent with the contrast described above: XDOAS fits the across-track structures out of the residual within the fit, while the S5P-PAL RMS field still contains the structures that its destriping corrects only afterwards, in the retrieved SCD. The retrieval uncertainty maps (bottom row) show the same pattern. The XDOAS uncertainty field is smooth and even, while the S5P-PAL field is noisy at high latitudes, and

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

this noise carries through to uncertainty over the whole orbit. The difference map (bottom-right) is small but follows a spatial pattern, which shows that the regularised covariance of XDOAS removes the pixel-to-pixel uncertainty fluctuations that the S5P-PAL product passes on to its uncertainty estimate.

Together, the maps confirm what the previous comparisons have shown XDOAS agrees with the S5P-PAL product across the whole orbit and over the polar vortex, while giving a smoother RMS residual and retrieval uncertainty field, most of all where the illumination is low and the OClO signal is measured.

3.2.5 Consistency of the reported uncertainty

The comparisons above establish that XDOAS and S5P-PAL recover the same geophysical signal. A separate and complementary question is whether the uncertainty that each product reports is consistent with the actual scatter of its slant columns. This is the calibration of the retrieval uncertainty, and it is independent of the agreement between the two products: two retrievals can recover the same columns while characterising their uncertainty differently.

To assess this without relying on the internal definition of the residual of either product, we use the observed pixel-to-pixel scatter of the slant columns over the equatorial clean sector ($-30°$ to $+30°$) as a neutral reference. In this sector the OClO column is close to zero, and the scatter of the retrieved SCDs is governed by noise alone and provides a direct, product-independent measure of the true retrieval scatter. For each product we compute the reduced chi-square

$$\chi^2_{\mathrm{red}} = \left\langle \frac{\left(S - \widetilde{S}\right)^2}{\sigma^2} \right\rangle, \tag{27}$$

where S is the retrieved SCD, $\widetilde{S}$ its median over the sector, and σ the reported retrieval uncertainty. A value $\chi^2_{\mathrm{red}} \approx 1$ indicates that the reported uncertainty matches the observed scatter, $\chi^2_{\mathrm{red}} > 1$ indicates that the uncertainty is underestimated, i.e. the columns scatter more than their reported error bars allow. The same clean-sector mask is applied to each product over its own pixels, and the metric compares each reported uncertainty against the scatter of the same geophysical region.

Figure 12 shows the consistency of the retrieval uncertainty over the equatorial clean sector. For XDOAS the reduced chi-square stays close to unity across all latitude bins (between about 0.97 and 1.02). The obtained uncertainty therefore follows the observed scatter throughout the sector. The left panel shows the same behaviour directly: the observed scatter and the obtained uncertainty of XDOAS overlap at all latitudes. The S5P-PAL product, in contrast, has a reduced chi-square systematically above unity (between about 1.1 and 1.45), and its reported uncertainty lies below the observed scatter throughout the sector. This corresponds to an underestimation of the retrieval uncertainty by roughly 10 to 20%, increasing towards the higher latitudes of the sector, where the solar zenith angle grows and the spectral fit quality begins to degrade. In absolute terms the observed scatter of S5P-PAL is in fact somewhat lower than that of XDOAS across most of the sector (left panel, solid curves), making S5P-PAL marginally more precise there; this does not affect the calibration result, since the reduced chi-square tests whether the reported uncertainty of each retrieval matches its own observed scatter, not which retrieval scatters less.

This confirms the consistency test of Sect. 3.1.5: the XDOAS uncertainty tracks the residual scatter of each pixel, while the S5P-PAL uncertainty does not increase enough where the scatter grows. This result supports the interpretation of the high-

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

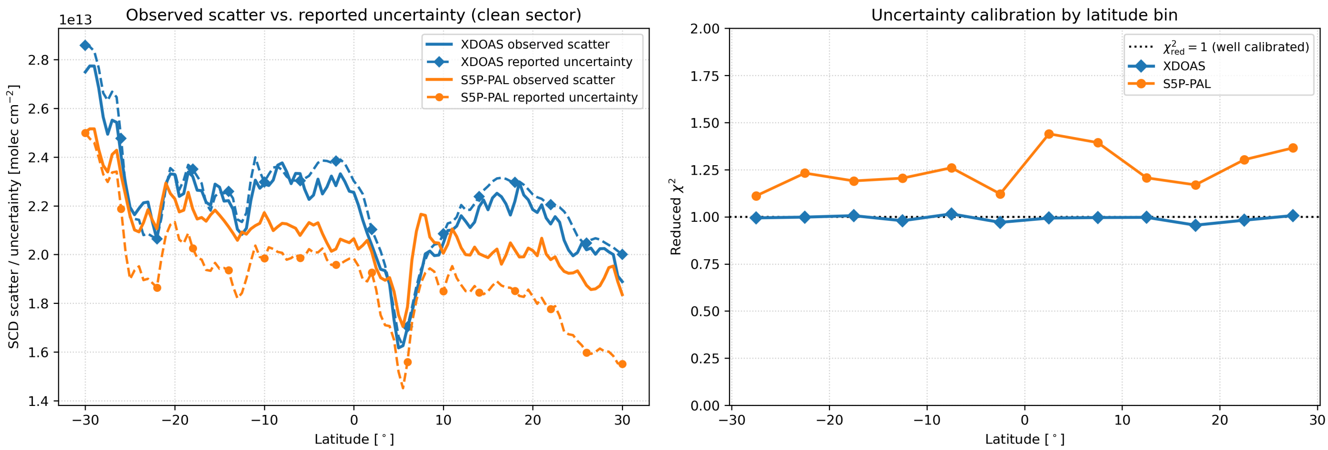

Figure 12. Consistency of the retrieval uncertainty over the equatorial clean sector. Left: running 1σ observed scatter of the slant columns (solid) and median reported uncertainty (dashed) as a function of latitude, for XDOAS (blue) and S5P-PAL (orange). Right: reduced chi-square in 5° latitude bins, the dotted line marks $\chi^2_{\text{red}} = 1$ (reported uncertainty equal to observed scatter).

latitude behaviour discussed in Sect. 3.2.4, where the wider XDOAS uncertainty was shown to reflect the larger residual scatter rather than a loss of precision. This metric characterises the consistency of the reported uncertainty rather than the accuracy of the retrieved columns. The latter was shown to agree between the two products in the previous comparisons. A full evaluation of its temporal stability is beyond the scope of the present study and will be addressed in future work.

4 Conclusions

The treatment of the fit residual is the central problem of classical DOAS retrievals. At the precision of modern satellite instruments such as TROPOMI, the residual of a classical fit is no longer an approximately white background: it splits into systematic structures of physical and instrumental origin and a correlated random noise. The classical fit treats both in the same way, which is its main limitation. Two approaches have been proposed in the recent DOAS literature, each addressing the problem from a different angle: covariance-weighted DOAS reweights the spectral fit using the residual covariance matrix of a reference ensemble, while PCA-DOAS adds the leading PCA components of that same ensemble to the fit matrix, which fit the systematic residual structures out of the residual explicitly. Each approach corrects part of the limitation of classical DOAS, but each shows a weakness when used alone, as this study shows.

XDOAS combines covariance weighting and PCA-based residual correction within a single consistent linear formulation. The central idea is the regularised covariance of Eq. (25), which replaces the variance along the PCA directions by a common background-noise level and reweights only the remaining correlated noise. The systematic and random parts of the residual are addressed by separate mechanisms acting on separate parts of the residual: the PCA components fit the systematic structures out of the residual, while the regularised covariance down-weights the remaining correlated noise. Classical DOAS, covariance-

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

weighted DOAS, and PCA-DOAS are singular cases of this more general formulation (Table 1, Sect. 2.5), which makes XDOAS a unifying framework rather than a fourth alternative.

The evaluation of the OClO retrievals against the three reference retrievals of the DOAS family (Sect. 3.1), performed on TROPOMI/S5P data with identical fit parameters for the four algorithms, gives a consistent picture of the advantages of the unified retrieval. The precision of XDOAS is comparable to that of covariance-weighted DOAS and PCA-DOAS in the equatorial clean sector and improves by about 20% over classical DOAS (Sect. 3.1.1). The retrieval-uncertainty distribution is narrowed by about 20% in FWHM relative to classical DOAS, with a FWHM essentially identical to PCA-DOAS and slightly higher than that of covariance-weighted DOAS, whose lower FWHM does not correspond to a real improvement (Sect. 3.1.5). The across-track striping of classical DOAS is removed to within $\pm 1 \times 10^{13}\,\mathrm{molec\,cm^{-2}}$ by the three advanced retrievals, and within this group XDOAS and PCA-DOAS reach the lowest RMS residual across the swath (Sect. 3.1.3). The spatial fields of SCD, RMS residual, and retrieval uncertainty are the smoothest of the family for XDOAS (Sect. 3.1.4).

The test that separates the four retrievals is the consistency between the RMS residual and the retrieval uncertainty (Sect. 3.1.5), evaluated over the equatorial clean sector. Classical DOAS shows an exact linear correlation ($R^2 = 1.000$), which is not a sign of quality but a trivial consequence of computing the uncertainty from the same unweighted residual that defines the RMS. Covariance-weighted DOAS loses this consistency completely: the points collapse into a nearly flat band ($R^2 = 0.002$, slope falling to 4.75×10^{14}), because the covariance weighting lowers the uncertainty of every pixel to a similar value regardless of its actual residual, which exposes the lower FWHM of Sect. 3.1.2 as an artefact rather than a real improvement. PCA-DOAS recovers an almost perfect correlation ($R^2 = 0.991$), with hardly any pixel-to-pixel dispersion, because the PCA components fit out part of the random noise together with the systematic structures and the uncertainty becomes fixed to the RMS residual. XDOAS provides the most balanced relationship between residual magnitude and retrieval uncertainty: a clean linear correlation with a moderate pixel-to-pixel spread ($R^2 = 0.977$, below that of PCA-DOAS). The retrieval uncertainty of XDOAS is neither compressed to a constant value nor fixed to the RMS residual, but follows the spectral fit quality of each pixel. Each of the reference methods improves one aspect of classical DOAS but has a weakness when used alone: covariance weighting lowers the measurement noise but decouples the uncertainty from the residual, and the PCA components fit out the systematic structures but also fit out part of the random noise, fixing the uncertainty to the residual. XDOAS is the only retrieval that combines the strengths of both, the lower noise of covariance weighting and the explicit correction of systematic structures of PCA, without taking on the weakness of either, because it applies a different method to each part of the residual. This is the advantage of XDOAS over the other three retrievals.

The comparison against the S5P-PAL OClO product (Sect. 3.2) confirms that these advantages carry over to a product that agrees with an independently developed reference. XDOAS is run with the same TROPOMI input data, the same spectral fitting window, the same baseline polynomial, and the same cross sections as the S5P-PAL product, and both products are filtered in the same way, so that they differ only in how the systematic residual structures are treated: within the fit in XDOAS, through the PCA pseudo-absorbers and the regularised covariance, and through empirical pseudo-absorbers and an explicit destriping after the fit in S5P-PAL. The two products agree closely ($R^2 = 0.922$, slope 1.065, median difference $1.05 \times 10^{12}\,\mathrm{molec\,cm^{-2}}$, about one order of magnitude below the typical retrieval uncertainty of either product), and the latitudinal comparison (Sect. 3.2.2)

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

shows that this small difference is uniform across the background and that the strong polar-vortex signal is reproduced by the two products within their respective uncertainty bands. The two products are therefore comparable in the columns themselves; over the clean sector S5P-PAL even shows slightly lower scatter, and XDOAS does not claim an improvement in SCD accuracy. Its advantages appear in the other diagnostics. The reported uncertainty of XDOAS is consistent with the observed clean-sector scatter (reduced chi-square close to unity), whereas the S5P-PAL uncertainty underestimates it by roughly 10 to 20% (reduced chi-square between about 1.1 and 1.45, Sect. 3.2.5). The RMS residual of XDOAS is lower across the swath and its uncertainty field is smoother, and under extreme optical paths (Sect. 3.2.3) this RMS gain grows with SZA, with the largest improvement near the day-night boundary where polar chlorine activation is observed. All of this follows from how the residual is treated: XDOAS handles the systematic structures and the random noise within the fit, through a basis derived from the measured residuals of each orbit, rather than through the fixed empirical corrections and destriping that S5P-PAL applies afterwards. Because this basis adapts to the conditions of each acquisition, independently of the target absorption, the retrieved SCD, the RMS residual, and the uncertainty remain consistent with one another. The present comparison is based on a single TROPOMI orbit, chosen to illustrate the behaviour of the four retrievals and of XDOAS against S5P-PAL under representative conditions; a multi-orbit and multi-season evaluation of the temporal stability of the product is the subject of ongoing work.

OClO is the most reliable optical tracer of activated chlorine in the polar stratosphere, and its retrieval from satellite is limited to the low-illumination conditions of polar winter and early spring, where the residual of a DOAS fit is most affected by systematic instrumental structures. By fitting these structures out of the residual through the PCA components and recovering an uncertainty that follows the spectral fit quality through the regularised covariance, XDOAS lowers the RMS residual exactly in this range and gives an uncertainty that can be trusted pixel by pixel. This is useful for OClO, where the signal is weak and the conditions are difficult, and where a reliable uncertainty helps to separate a real chlorine activation signal from the residual scatter. The comparison with the S5P-PAL product shows that these gains are obtained without applying empirical post-processing corrections. XDOAS agrees with the S5P-PAL product across the orbit and over the polar vortex, while showing a lower RMS residual and a reported uncertainty that remains consistent with the observed scatter of the slant columns, where the S5P-PAL uncertainty is underestimated.

XDOAS is the baseline algorithm for the operational OClO product of the Copernicus Sentinel-5 mission, within the Atmospheric Composition Satellite Application Facility (AC-SAF) of EUMETSAT. Because XDOAS fits the systematic residual structures out of the residual explicitly and recovers an uncertainty that follows the spectral fit quality of each pixel, it is well suited to deliver a stable and well characterised OClO record over the lifetime of the mission. Another practical advantage of XDOAS is that, with an appropriately chosen reference region for the covariance and PCA components, the dominant across-track structures are fitted out within the retrieval itself, removing the need for a separate destriping step. This simplifies the implementation in operational ground segments.

Although XDOAS has been applied here for OClO, the method is general and can be applied to any trace gas with narrow absorption bands in the ultraviolet or visible. XDOAS works on the fit residual and not on the specific absorber: the systematic structures are fitted out of the residual through the PCA components, and the remaining correlated noise is down-weighted by the regularised covariance, independently of which gas is being retrieved. This makes XDOAS especially well suited to

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

extreme viewing geometries typical in geostationary sounders, such as Sentinel-4, GEMS, and TEMPO, whose dense temporal sampling over the mid-latitudes is particularly valuable for species with a strong tropospheric signal such as NO_2 and volatile organic compounds (e.g. CHOCHO and HCHO). These applications are the subject of ongoing work.

Appendix A: Statistical derivation of the XDOAS estimators

This appendix develops the statistical reasoning behind the four retrievals introduced in Sect. 2. The main text presents the estimators in the form actually used by the algorithm; here we make explicit the error model, the assumptions, and the steps that connect classical least squares to the regularised covariance-weighted fit of XDOAS. None of this is required to apply the method, but it clarifies why each estimator takes the form it does and under which conditions it is valid.

A1 Data model and the two error contributions

The starting point is the linearised data model of Eq. (4),

$$\boldsymbol{\tau}^{\delta} = \mathbf{B}\boldsymbol{\beta}^{\dagger} + \boldsymbol{\delta}, \tag{A1}$$

where $\boldsymbol{\beta}^{\dagger}$ is the (unknown) true state vector and $\boldsymbol{\delta} \in \mathbb{R}^{M}$ is the data error, i.e. everything in the measured optical depth that the forward model $\mathbf{B}\boldsymbol{\beta}$ does not represent. Following Eq. (15), the data error is split into a systematic and a random part,

$$\boldsymbol{\delta} = \boldsymbol{\delta}_{\mathrm{sys}} + \boldsymbol{\delta}_{\mathrm{rnd}}, \qquad \mathbb{E}[\boldsymbol{\delta}_{\mathrm{rnd}}] = \mathbf{0}, \qquad \mathbb{E}[\boldsymbol{\delta}] = \boldsymbol{\delta}_{\mathrm{sys}}. \tag{A2}$$

The systematic part $\boldsymbol{\delta}_{\mathrm{sys}}$ collects reproducible structures of physical or instrumental origin (residual stray light, ISRF mismatch, across-track artefacts, undersampling), which recur with a stable spectral shape across many observations. The random part $\boldsymbol{\delta}_{\mathrm{rnd}}$ is zero-mean but, for a hyperspectral instrument, is not spectrally white: it is correlated over several detector pixels, and is characterised by its covariance $\mathrm{Cov}(\boldsymbol{\delta}_{\mathrm{rnd}}) = \mathbb{E}[\boldsymbol{\delta}_{\mathrm{rnd}}\boldsymbol{\delta}_{\mathrm{rnd}}^{\top}]$. The four retrievals differ only in how they treat these two contributions.

A2 From ordinary to generalized least squares

The ordinary least-squares (OLS) estimator of Eq. (6) minimises the unweighted sum of squared residuals, $\frac{1}{2}\|\mathbf{B}\boldsymbol{\beta} - \boldsymbol{\tau}^{\delta}\|_2^2$. This is the maximum-likelihood estimator when the data error is white, i.e. when $\mathrm{Cov}(\boldsymbol{\delta}) = \sigma^2\mathbf{I}_M$. When instead the data error has a known covariance $\mathbf{C}_{\delta}$, the Gauss–Markov theorem states that the best linear unbiased estimator (BLUE) is the generalized least-squares (GLS) solution, obtained by minimising the Mahalanobis distance

$$\mathcal{F}_{\mathrm{GLS}}(\boldsymbol{\beta}) = \tfrac{1}{2}\left(\mathbf{B}\boldsymbol{\beta} - \boldsymbol{\tau}^{\delta}\right)^{\top}\mathbf{C}_{\delta}^{-1}\left(\mathbf{B}\boldsymbol{\beta} - \boldsymbol{\tau}^{\delta}\right), \tag{A3}$$

whose stationarity condition $\partial\mathcal{F}_{\mathrm{GLS}}/\partial\boldsymbol{\beta} = \mathbf{0}$ gives the normal equations $\mathbf{B}^{\top}\mathbf{C}_{\delta}^{-1}\mathbf{B}\hat{\boldsymbol{\beta}} = \mathbf{B}^{\top}\mathbf{C}_{\delta}^{-1}\boldsymbol{\tau}^{\delta}$ and hence Eq. (14). The OLS estimator is the special case $\mathbf{C}_{\delta} = \mathbf{I}_M$: uniform weighting assumes, implicitly, that all spectral directions carry independent errors of equal variance. For a modern hyperspectral instrument this assumption is not met, which is the formal reason classical DOAS leaves correlated residual structure unexploited.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

A3 The residual as a proxy for the data error

In practice $\mathbf{C}_\delta$ is unknown and must be estimated. The only quantity available after a fit is the residual of Eq. (7), $\mathbf{r} = \boldsymbol{\tau}^\delta - \mathbf{B}\hat{\boldsymbol{\beta}}$. Substituting the data model (A1) gives the exact relation of Eq. (8),

$$\mathbf{r} = \boldsymbol{\delta} + \mathbf{B}\left(\boldsymbol{\beta}^\dagger - \hat{\boldsymbol{\beta}}\right). \tag{A4}$$

The residual therefore equals the data error plus a contribution from the fit error $\boldsymbol{\beta}^\dagger - \hat{\boldsymbol{\beta}}$ projected through $\mathbf{B}$. The central approximation enabling the empirical approach is

(A1) *Accurate pre-fit.* If the preliminary fit removes the modelled absorber and broadband signals well, the fit error is small, $\hat{\boldsymbol{\beta}} \approx \boldsymbol{\beta}^\dagger$, so that $\mathbf{r} \approx \boldsymbol{\delta}$ and the statistics of the residual ensemble reproduce those of the data error.

Under (A1), an ensemble of N_{ref} residuals $\mathbf{r}_n$ collected over a clean reference scene yields the empirical estimates of Eq. (16): the mean residual $\bar{\mathbf{r}}$ and the covariance $\mathbf{C}$. Two further assumptions are needed for these estimates to be meaningful for an arbitrary retrieval pixel:

(A2) *Representativeness / stationarity.* The error statistics of the reference ensemble are sufficiently similar to those of the spectra to which the correction is applied; i.e. $\bar{\mathbf{r}}$ and $\mathbf{C}$ are approximately stationary between the ensemble and the retrieval.

(A3) *Target-gas exclusion.* The target absorption is not significantly present in the reference ensemble, so that neither $\bar{\mathbf{r}}$ nor $\mathbf{C}$ contains the signal of interest. Otherwise the correction would remove part of the true column (target-signal absorption); this motivates the clean-sector pre-fit and, for strong-background absorbers, the adaptive step of Sect. 2.4.1.

Because the mean is subtracted in Eq. (16), $\mathbf{C}$ estimates the covariance of the *fluctuating* part of the residual. Combining this with the decomposition (A2) gives the identifications used throughout the main text,

$$\mathbf{r}_{\mathrm{sys}} \approx \bar{\mathbf{r}}, \qquad \mathrm{Cov}(\boldsymbol{\delta}_{\mathrm{rnd}}) \approx \mathbf{C}. \tag{A5}$$

A4 Covariance-weighted DOAS as GLS with the estimated covariance

Covariance-weighted DOAS applies the GLS estimator (A3) with the unknown $\mathbf{C}_\delta$ replaced by its empirical estimate $\mathbf{C}$. The systematic structure is not down-weighted but *modelled*: the mean residual $\bar{\mathbf{r}}$ is appended to the fit matrix as a pseudo-absorber (Eq. 17). Its amplitude is left free through a coefficient c_{R} because, while the *shape* of the systematic structure is reproducible (assumption A2), its *amplitude* varies from observation to observation; fixing it to the ensemble value would over- or under-correct individual spectra. The resulting estimator is Eq. (18).

A5 PCA-based correction and the variance interpretation

The PCA-based correction instead represents the systematic structure explicitly through the leading eigenvectors of $\mathbf{C}$. Since $\mathbf{C}$ is symmetric and positive semi-definite, it admits the orthonormal eigendecomposition of Eq. (19), with $\mathbf{v}_i^\top \mathbf{v}_j = \delta_{ij}$. Ex-

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

panding a centred residual in this basis, $\mathbf{r}_n - \bar{\mathbf{r}} = \sum_k \alpha_{nk}\mathbf{v}_k$ with $\alpha_{nk} = \mathbf{v}_k^\top(\mathbf{r}_n - \bar{\mathbf{r}})$, and using orthonormality, the ensemble variance of the k-th coefficient is

$$\mathrm{Var}(\alpha_{\cdot k}) = \mathbf{v}_k^\top \mathbf{C}\,\mathbf{v}_k = \lambda_k. \tag{A6}$$

Each eigenvalue is therefore the variance of the residual ensemble along its eigenvector, and the leading K directions are those along which the residual varies most. These K eigenvectors are added to the fit as pseudo-absorbers (Eq. 20) and fitted by OLS. Two assumptions govern this step:

(A4) *Low-rank systematic structure.* The systematic part of the residual is well approximated by the span of the first K eigenvectors, $\mathbf{r}_{\mathrm{sys}} \approx \sum_{k=1}^{K} c_k \mathbf{v}_k$, with $K \ll M$.

(A5) *Choice of K.* K is chosen large enough to capture the dominant structures but small enough to avoid absorbing physical signal (overfitting). In practice K is a fixed configuration parameter of the retrieval, with the eigenvalue spectrum $\lambda_1 \geq \cdots \geq \lambda_M$ providing guidance on how many components carry significant residual variance. For the OClO retrieval, $K = 4$ is used (Sect. 2.4); increasing K beyond this value produced no appreciable further reduction of the residual or change in the retrieved SCDs.

PCA-DOAS is fitted by OLS rather than GLS by design: it isolates the effect of modelling the systematic structures, leaving the weighting of the random part to the covariance-based retrievals. Weighting the PCA fit with the full $\mathbf{C}$ would suppress the explicitly fitted directions a second time (see Sect. A6), which is precisely the redundancy that the regularised XDOAS covariance is constructed to avoid.

A6 Why the XDOAS covariance must be regularised

XDOAS combines both treatments: the systematic structure is modelled by the mean residual and the leading K eigenvectors (Eq. 23), while the remaining random error is down-weighted by a covariance. The natural weighting matrix would be the covariance of the part of the residual *not* already in the fit matrix,

$$\mathrm{Cov}(\boldsymbol{\delta}_{\mathrm{rnd}})\Big|_{k>K} = \sum_{k=K+1}^{M} \lambda_k\,\mathbf{v}_k\mathbf{v}_k^\top. \tag{A7}$$

This matrix has rank $M-K$: it is identically zero along $\mathbf{v}_1, \ldots, \mathbf{v}_K$ and therefore *singular* whenever $K > 0$. Its inverse does not exist, so it cannot be used directly as a GLS weighting matrix. The retrieval resolves this by replacing the variance along the first K directions with a common background level rather than zero,

$$\sigma^2_{\mathrm{noise}} = \frac{1}{M-K}\sum_{k=K+1}^{M} \lambda_k, \tag{A8}$$

yielding the regularised covariance $\mathbf{C}_{\mathrm{reg}}$ and its inverse $\mathbf{C}_{\mathrm{reg}}^{-1}$ of Eq. (25). Two remarks make the nature of this step precise:

Table A1. The reference retrievals as limiting cases of the XDOAS estimator of Eq. (26). K is the number of PCA pseudo-absorbers and $\mathbf{W}$ the weighting matrix ($\mathbf{C}_{\mathrm{reg}}$ with the stated choice).

Retrieval	K	$\mathbf{W}$	Treatment
Classical DOAS	0	$\mathbf{I}_M$	none
Covariance-weighted	0	$\mathbf{C}$	random part only
PCA-DOAS	> 0	$\mathbf{I}_M$	systematic part only
XDOAS	> 0	$\mathbf{C}_{\mathrm{reg}}$	both, no double counting

(A6) *Modelling assumption, not estimation.* The replacement $\lambda_k \to \sigma^2_{\mathrm{noise}}$ for $k \leq K$ is a modelling choice. It is not implied by estimation theory; its role is to keep $\mathbf{C}_{\mathrm{reg}}$ positive-definite and invertible while assigning to the explicitly fitted directions only the residual background-noise variance, rather than the large variances λ_k that those directions carry in $\mathbf{C}$.

(A7) *Avoidance of double suppression.* Leaving the original large λ_k ($k \leq K$) in the weighting matrix would penalise the very directions already represented in $\mathbf{B}_{\mathrm{X}}$, suppressing them once through the fit and once through the weighting. Flattening these eigenvalues to $\sigma^2_{\mathrm{noise}}$ removes that redundancy.

With $\mathbf{C}^{-1}_{\mathrm{reg}}$ defined, the XDOAS estimate is the GLS solution of Eq. (26). The construction guarantees that the two mechanisms act on disjoint parts of the residual: the pseudo-absorbers in $\mathbf{B}_{\mathrm{X}}$ remove the systematic structures (directions $k \leq K$ plus the mean), while $\mathbf{C}^{-1}_{\mathrm{reg}}$ shapes the weighting of the correlated random noise that remains (directions $k > K$).

A7 The reference methods as limiting cases

The estimator (26) reproduces the three reference retrievals for particular choices of K and of the weighting matrix, summarised in Table A1. This is the formal statement of the unifying property discussed in Sect. 2.5: classical, covariance-weighted, and PCA-DOAS are not separate algorithms but special cases of a single regularised GLS fit.

Code and data availability. The XDOAS code can be requested from the authors. The TROPOMI L1 data used in this study are available from https://dataspace.copernicus.eu/ (last access: 13 July 2026). The TROPOMI S5P-PAL OClO data are available from https://data-portal.s5p-pal.com/ (last access: 13 July 2026).

Author contributions. LA developed the XDOAS framework and wrote the paper. SS tested the XDOAS framework and applied it to the retrieval of NO_2. AD contributed to Sect. 2 and Appendix A. AR participated in the comparison with the TROPOMI S5P-PAL OClO data. DL contributed to the text and organized the funding. All authors contributed to the discussion of the results and to the editing of the manuscript.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Competing interests. At least one of the (co-)authors is a member of the editorial board of Atmospheric Measurement Techniques. The authors have no other competing interests to declare.

Acknowledgements. The S5P-PAL OClO product used for the comparison was developed at the Institute of Environmental Physics (IUP), University of Bremen, within the framework of the Sentinel-5p+ Innovation (S5p+I) activity, and was made available through the S5P-PAL data portal. During the preparation of this manuscript, the authors used an AI-based language assistant (Claude, Anthropic) to improve the clarity and wording of the English text. The tool was used for language editing only and did not contribute to the scientific content, the data analysis, or the conclusions, which are entirely the responsibility of the authors.

Financial support. The development of the XDOAS framework was funded by the German Aerospace Center (DLR KTR 2472044), and its application to OClO was funded by the EUMETSAT/DLR AC-SAF project.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

References

Alvarado, L. M. A., Richter, A., Vrekoussis, M., Wittrock, F., Hilboll, A., Schreier, S. F., and Burrows, J. P.: An improved glyoxal retrieval from OMI measurements, Atmospheric Measurement Techniques, 7, 4133–4150, https://doi.org/10.5194/amt-7-4133-2014, 2014.

Alvarado, L. M. A., Richter, A., Vrekoussis, M., Hilboll, A., Kalisz Hedegaard, A. B., Schneising, O., and Burrows, J. P.: Unexpected long-range transport of glyoxal and formaldehyde observed from the Copernicus Sentinel-5 Precursor satellite during the 2018 Canadian wildfires, Atmospheric Chemistry and Physics, 20, 2057–2072, https://doi.org/10.5194/acp-20-2057-2020, 2020.

Beirle, S., Lampel, J., Lerot, C., Sihler, H., and Wagner, T.: Linearisation of the effects of spectral shift and stretch in DOAS analysis, Atmospheric Measurement Techniques, 6, 661–675, https://doi.org/10.5194/amt-6-661-2013, 2013.

Bogumil, K., Orphal, J., Homann, T., Voigt, S., Spietz, P., Fleischmann, O. C., Vogel, A., Hartmann, M., Kromminga, H., Bovensmann, H., Frerick, J., and Burrows, J. P.: Measurements of molecular absorption spectra with the SCIAMACHY pre-flight model: instrument characterization and reference data for atmospheric remote sensing in the 230–2380 nm spectral region, Journal of Photochemistry and Photobiology A: Chemistry, 157, 167–184, https://doi.org/10.1016/S1010-6030(03)00062-5, 2003.

Bösch, T., Rozanov, V., Richter, A., Peters, E., Rozanov, A., Wittrock, F., Merlaud, A., Lampel, J., Schmitt, S., de Haij, M., Berkhout, S., Henzing, B., Apituley, A., den Hoed, M., Vonk, J., Tiefengraber, M., Müller, M., and Burrows, J. P.: BOREAS – a new MAX-DOAS profile retrieval algorithm for aerosols and trace gases, Atmospheric Measurement Techniques, 11, 6833–6859, https://doi.org/10.5194/amt-11-6833-2018, 2018.

Bramstedt, K., Gleason, J., Loyola, D., Thomas, W., Bracher, A., Weber, M., and Burrows, J. P.: Comparison of total ozone from the satellite instruments GOME and TOMS with measurements from the Dobson network 1996–2000, Atmospheric Chemistry and Physics, 3, 1409–1419, https://doi.org/10.5194/acp-3-1409-2003, 2003.

Burrows, J. P., Weber, M., Buchwitz, M., Rozanov, V., Ladstätter-Weißenmayer, A., Richter, A., DeBeek, R., Hoogen, R., Bramstedt, K., Eichmann, K.-U., Eisinger, M., and Perner, D.: The Global Ozone Monitoring Experiment (GOME): Mission concept and first scientific results, Journal of the Atmospheric Sciences, 56, 151–175, https://doi.org/10.1175/1520-0469(1999)056<0151:TGOMEG>2.0.CO;2, 1999.

Chance, K. and Kurucz, R. L.: An improved high-resolution solar reference spectrum for earth's atmosphere measurements in the ultraviolet, visible, and near infrared, Journal of Quantitative Spectroscopy and Radiative Transfer, 111, 1289–1295, https://doi.org/10.1016/j.jqsrt.2010.01.036, 2010.

Coldewey-Egbers, M., Weber, M., Buchwitz, M., and Burrows, J. P.: Application of a modified DOAS method for total ozone retrieval from GOME data at high polar latitudes, Advances in Space Research, 35, 749–753, https://doi.org/10.1016/j.asr.2005.03.122, 2005.

Coldewey-Egbers, M., Loyola, D. G., Lerot, C., and Van Roozendael, M.: Global, regional and seasonal analysis of total ozone trends derived from the 1995–2020 GTO-ECV climate data record, Atmospheric Chemistry and Physics, 22, 6861–6878, https://doi.org/10.5194/acp-22-6861-2022, 2022.

De Smedt, I., Theys, N., Yu, H., Danckaert, T., Lerot, C., Compernolle, S., Van Roozendael, M., Richter, A., Hilboll, A., Peters, E., Pedergnana, M., Loyola, D., Beirle, S., Wagner, T., Eskes, H., van Geffen, J., Boersma, K. F., and Veefkind, P.: Algorithm theoretical baseline for formaldehyde retrievals from S5P TROPOMI and from the QA4ECV project, Atmospheric Measurement Techniques, 11, 2395–2426, https://doi.org/10.5194/amt-11-2395-2018, 2018.

Farman, J. C., Gardiner, B. G., and Shanklin, J. D.: Large losses of total ozone in Antarctica reveal seasonal ClOx/NOx interaction, Nature, 315, 207–210, https://doi.org/10.1038/315207a0, 1985.

Garane, K., Koukouli, M.-E., Verhoelst, T., Lerot, C., Heue, K.-P., Fioletov, V., Balis, D., Bais, A., Bazureau, A., Dehn, A., Goutail, F., Granville, J., Griffin, D., Hubert, D., Keppens, A., Lambert, J.-C., Loyola, D., McLinden, C., Pazmino, A., Pommereau, J.-P., Redondas, A., Romahn, F., Valks, P., Van Roozendael, M., Xu, J., Zehner, C., Zerefos, C., and Zimmer, W.: TROPOMI/S5P total ozone column data: global ground-based validation and consistency with other satellite missions, Atmospheric Measurement Techniques, 12, 5263–5287, https://doi.org/10.5194/amt-12-5263-2019, 2019.

González Abad, G., Liu, X., Chance, K., Wang, H., Kurosu, T. P., and Suleiman, R.: Updated Smithsonian Astrophysical Observatory Ozone Monitoring Instrument (SAO OMI) formaldehyde retrieval, Atmospheric Measurement Techniques, 8, 19–32, https://doi.org/10.5194/amt-8-19-2015, 2015.

Hönninger, G., von Friedeburg, C., and Platt, U.: Multi axis differential optical absorption spectroscopy (MAX-DOAS), Atmospheric Chemistry and Physics, 4, 231–254, https://doi.org/10.5194/acp-4-231-2004, 2004.

Höpfner, M., Deshler, T., Pitts, M., Poole, L., Spang, R., Stiller, G., and von Clarmann, T.: The MIPAS/Envisat climatology (2002–2012) of polar stratospheric cloud volume density profiles, Atmospheric Measurement Techniques, 11, 5901–5923, https://doi.org/10.5194/amt-11-5901-2018, 2018.

Joiner, J., Marchenko, S., Fasnacht, Z., Lamsal, L., Li, C., Vasilkov, A., and Krotkov, N.: Use of machine learning and principal component analysis to retrieve nitrogen dioxide (NO2) with hyperspectral imagers and reduce noise in spectral fitting, Atmospheric Measurement Techniques, 16, 481–500, https://doi.org/10.5194/amt-16-481-2023, 2023.

Lerot, C., Hendrick, F., Van Roozendael, M., Alvarado, L. M. A., Richter, A., De Smedt, I., Theys, N., Vlietinck, J., Yu, H., Van Gent, J., Stavrakou, T., Müller, J.-F., Valks, P., Loyola, D., Irie, H., Kumar, V., Wagner, T., Schreier, S. F., Sinha, V., Wang, T., Wang, P., and Retscher, C.: Glyoxal tropospheric column retrievals from TROPOMI – multi-satellite intercomparison and ground-based validation, Atmospheric Measurement Techniques, 14, 7775–7807, https://doi.org/10.5194/amt-14-7775-2021, 2021.

Li, C., Joiner, J., Krotkov, N. A., and Bhartia, P. K.: A fast and sensitive new satellite SO2 retrieval algorithm based on principal component analysis: Application to the ozone monitoring instrument, Geophysical Research Letters, 40, 6314–6318, https://doi.org/10.1002/2013GL058134, 2013.

Liu, S., Valks, P., Pinardi, G., Xu, J., Chan, K. L., Argyrouli, A., Lutz, R., Beirle, S., Khorsandi, E., Baier, F., Huijnen, V., Bais, A., Donner, S., Dörner, S., Gratsea, M., Hendrick, F., Karagkiozidis, D., Lange, K., Piters, A. J. M., Remmers, J., Richter, A., Van Roozendael, M., Wagner, T., Wenig, M., and Loyola, D. G.: An improved TROPOMI tropospheric NO_2 research product over Europe, Atmospheric Measurement Techniques, 14, 7297–7327, https://doi.org/10.5194/amt-14-7297-2021, 2021.

Loyola, D. G., Gimeno García, S., Lutz, R., Argyrouli, A., Romahn, F., Spurr, R. J. D., Pedergnana, M., Doicu, A., Molina García, V., and Schüssler, O.: The operational cloud retrieval algorithms from TROPOMI on board Sentinel-5 Precursor, Atmospheric Measurement Techniques, 11, 409–427, https://doi.org/10.5194/amt-11-409-2018, 2018.

McElroy, M. B., Salawitch, R. J., Wofsy, S. C., and Logan, J. A.: Reductions of Antarctic ozone due to synergistic interactions of chlorine and bromine, Nature, 321, 759–762, https://doi.org/10.1038/321759a0, 1986.

Meier, A., Richter, A., Pinardi, G., and Lerot, C.: Monitoring Stratospheric OClO with Sentinel-5p: Algorithm Theoretical Baseline Document, Algorithm Theoretical Baseline Document S5p+I_OCLO_IUP-UB_ATBD, Institute of Environmental Physics (IUP), University of Bremen, sentinel-5p + Innovation, Theme 2: Chlorine Dioxide. Available at https://data-portal.s5p-pal.com/product-docs/oclo/S5p+I_OClO_IUP-UB_ATBD_v3.1.pdf, 2024.

Noxon, J. F.: Nitrogen dioxide in the stratosphere and troposphere measured by ground-based absorption spectroscopy, Science, 189, 547–549, https://doi.org/10.1126/science.189.4202.547, 1975.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Oetjen, H., Wittrock, F., Richter, A., Chipperfield, M. P., Medeke, T., Sheode, N., Sinnhuber, B.-M., Sinnhuber, M., and Burrows, J. P.: Evaluation of stratospheric chlorine chemistry for the Arctic spring 2005 using modelled and measured OClO column densities, Atmospheric Chemistry and Physics, 11, 689–703, https://doi.org/10.5194/acp-11-689-2011, 2011.

Palmer, P. I., Jacob, D. J., Chance, K., Martin, R. V., Spurr, R. J. D., Kurosu, T. P., Bey, I., Yantosca, R., Fiore, A., and Li, Q.: Air mass factor formulation for spectroscopic measurements from satellites: Application to formaldehyde retrievals from the Global Ozone Monitoring Experiment, Journal of Geophysical Research: Atmospheres, 106, 14 539–14 550, https://doi.org/10.1029/2000JD900772, 2001.

Peter, T. and Grooß, J.-U.: Polar stratospheric clouds and sulfate aerosol particles: microphysics, denitrification and heterogeneous chemistry, in: Stratospheric Ozone Depletion and Climate Change, edited by Müller, R., pp. 108–144, Royal Society of Chemistry, https://doi.org/10.1039/9781849733182-00108, 2012.

Platt, U. and Stutz, J.: Differential Optical Absorption Spectroscopy: Principles and Applications, Springer, Berlin, Heidelberg, https://doi.org/10.1007/978-3-540-75776-4, 2008.

Platt, U., Perner, D., and Pätz, H. W.: Simultaneous measurement of atmospheric CH2O, O3, and NO2 by differential optical absorption, Journal of Geophysical Research, 84, 6329–6335, https://doi.org/10.1029/JC084iC10p06329, 1979.

Puķīte, J., Kühl, S., Deutschmann, T., Platt, U., and Wagner, T.: Extending differential optical absorption spectroscopy for limb measurements in the UV, Atmospheric Measurement Techniques, 3, 631–653, https://doi.org/10.5194/amt-3-631-2010, 2010.

Puķīte, J., Borger, C., Dörner, S., Gu, M., Frieß, U., Meier, A. C., Enell, C.-F., Raffalski, U., Richter, A., and Wagner, T.: Retrieval algorithm for OClO from TROPOMI (TROPOspheric Monitoring Instrument) by differential optical absorption spectroscopy, Atmospheric Measurement Techniques, 14, 7595–7625, https://doi.org/10.5194/amt-14-7595-2021, 2021.

Puķīte, J., Borger, C., Dörner, S., Gu, M., and Wagner, T.: OClO as observed by TROPOMI: a comparison with meteorological parameters and polar stratospheric cloud observations, Atmospheric Chemistry and Physics, 22, 245–272, https://doi.org/10.5194/acp-22-245-2022, 2022.

Richter, A.: Personal communication, 11 May 2026, 2026.

Richter, A., Begoin, M., Hilboll, A., and Burrows, J. P.: An improved NO2 retrieval for the GOME-2 satellite instrument, Atmospheric Measurement Techniques, 4, 1147–1159, https://doi.org/10.5194/amt-4-1147-2011, 2011.

Rozanov, V. V. and Rozanov, A. V.: Differential optical absorption spectroscopy (DOAS) and air mass factor concept for a multiply scattering vertically inhomogeneous medium: theoretical consideration, Atmospheric Measurement Techniques, 3, 751–780, https://doi.org/10.5194/amt-3-751-2010, 2010.

Schiller, C. and Wahner, A.: Comment on “Stratospheric chlorine dioxide measurements with the AROTEL instrument”, Geophysical Research Letters, 23, 1531–1534, https://doi.org/10.1029/96GL01471, 1996.

Serdyuchenko, A., Gorshelev, V., Weber, M., Chehade, W., and Burrows, J. P.: High spectral resolution ozone absorption cross-sections — Part 2: Temperature dependence, Atmospheric Measurement Techniques, 7, 625–636, https://doi.org/10.5194/amt-7-625-2014, 2014.

Solomon, S.: Stratospheric ozone depletion: A review of concepts and history, Reviews of Geophysics, 37, 275–316, https://doi.org/10.1029/1999RG900008, 1999.

Solomon, S., Garcia, R. R., Rowland, F. S., and Wuebbles, D. J.: On the depletion of Antarctic ozone, Nature, 321, 755–758, https://doi.org/10.1038/321755a0, 1986.

Solomon, S., Mount, G. H., Sanders, R. W., and Schmeltekopf, A. L.: Visible spectroscopy at McMurdo Station, Antarctica: 2. Observations of OClO, Journal of Geophysical Research, 92, 8329–8338, https://doi.org/10.1029/JD092iD07p08329, 1987.

https://doi.org/10.5194/egusphere-2026-4360
Preprint. Discussion started: 31 August 2026

Thalman, R. and Volkamer, R.: Temperature dependent absorption cross-sections of O_2–O_2 collision pairs between 340 and 630 nm and at atmospherically relevant pressure, Physical Chemistry Chemical Physics, 15, 15 371–15 381, https://doi.org/10.1039/c3cp50968k, 2013.

Theys, N., Hedelt, P., De Smedt, I., Lerot, C., Yu, H., Vlietinck, J., Pedergnana, M., Arellano, S., Galle, B., Fernandez, D., Carlito, C. J. M., Barrington, C., Taisne, B., Delgado-Granados, H., Loyola, D., and Van Roozendael, M.: Sulfur dioxide retrievals from TROPOMI onboard Sentinel-5 Precursor: algorithm theoretical basis, Atmospheric Measurement Techniques, 10, 119–153, https://doi.org/10.5194/amt-10-119-2017, 2017.

Theys, N., Fioletov, V., Li, C., De Smedt, I., Lerot, C., McLinden, C., Krotkov, N., Griffin, D., Clarisse, L., Hedelt, P., Loyola, D., Wagner, T., Kumar, V., Innes, A., Ribas, R., Hendrick, F., Vlietinck, J., Brenot, H., and Van Roozendael, M.: A sulfur dioxide Covariance-Based Retrieval Algorithm (COBRA): application to TROPOMI reveals new emission sources, Atmospheric Chemistry and Physics, 21, 16 727–16 744, https://doi.org/10.5194/acp-21-16727-2021, 2021.

Theys, N., De Smedt, I., Danckaert, T., Richter, A., Vlietinck, J., and Van Roozendael, M.: Covariance-DOAS (CODOAS), the best of the two worlds?, Presentation at the International DOAS Workshop 2025, Thessaloniki, Greece, 9–11 July 2025, https://www.doas2025.gr/wp-content/uploads/2025/07/DOAS-2025-Programme_final-250g8w42.pdf, programme, Session 3: Instruments and algorithms; accessed: 12 May 2026, 2025.

van Geffen, J., Eskes, H., Compernolle, S., Pinardi, G., Verhoelst, T., Lambert, J.-C., Sneep, M., ter Linden, M., Ludewig, A., Boersma, K. F., and Veefkind, J. P.: Sentinel-5P TROPOMI NO2 retrieval: impact of version v2.2 improvements and comparisons with OMI and ground-based data, Atmospheric Measurement Techniques, 15, 2037–2060, https://doi.org/10.5194/amt-15-2037-2022, 2022.

Vandaele, A. C., Hermans, C., Simon, P. C., Carleer, M., Colin, R., Fally, S., Mérienne, M. F., Jenouvrier, A., and Coquart, B.: Measurements of the NO_2 absorption cross section from 42 000 cm^{-1} to 10 000 cm^{-1} (238–1000 nm) at 220 K and 294 K, Journal of Quantitative Spectroscopy and Radiative Transfer, 59, 171–184, https://doi.org/10.1016/S0022-4073(97)00168-4, 1998.

Veefkind, J. P., Aben, I., McMullan, K., Förster, H., de Vries, J., Otter, G., Claas, J., Eskes, H. J., de Haan, J. F., Kleipool, Q., van Weele, M., Hasekamp, O., Hoogeveen, R., Landgraf, J., Snel, R., Tol, P., Ingmann, P., Voors, R., Kruizinga, B., Vink, R., Visser, H., and Levelt, P. F.: TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications, Remote Sensing of Environment, 120, 70–83, https://doi.org/10.1016/j.rse.2011.09.027, 2012.

Wilmouth, D. M., Hanisco, T. F., Donahue, N. M., and Anderson, J. G.: Fourier transform ultraviolet spectroscopy of the $A\,^2\Pi_{3/2} \leftarrow X\,^2\Pi_{3/2}$ transition of BrO, The Journal of Physical Chemistry A, 103, 8935–8945, https://doi.org/10.1021/jp991651o, 1999.

World Meteorological Organization: Scientific Assessment of Ozone Depletion: 2022, Global Ozone Research and Monitoring Project, GAW Report No. 278, 509 pp., 2022.