The paper is interesting as a source of homogenized structural uncertainty studies, where the ROMEX database has been used to provide a coherent sampling of different source data. It is relevant to be aware of difficult to reduce errors, sometimes systematic, and that some atmospheric circumstances lead to retrievals of better quality, while others are quite dependent on apparently minor details of the retrieval. The issue of the retrieval uncertainty is still insufficiently addressed.

The subject is thus interesting and a good application of the ROMEX database. There however are a number of corrections and clarifications that I believe are important, as detailed below.

Reply:
Thank you for your constructive suggestions, which have helped improve the quality and clarity of this manuscript.

L150 and Tab1: The volume of Tianmu is indeed small within the ROMEX ensemble. However the average I have, admittedly on a superficial count, is of the order of 300 profiles/day, rather than 100/day. Please verify.

Reply:
Thank you for pointing this out. We verified and updated the daily profile count for Tianmu (~300 profiles/day) in Table 1 and the corresponding text (Line 150) to ensure the presented numbers are accurate. You can see the daily number below.

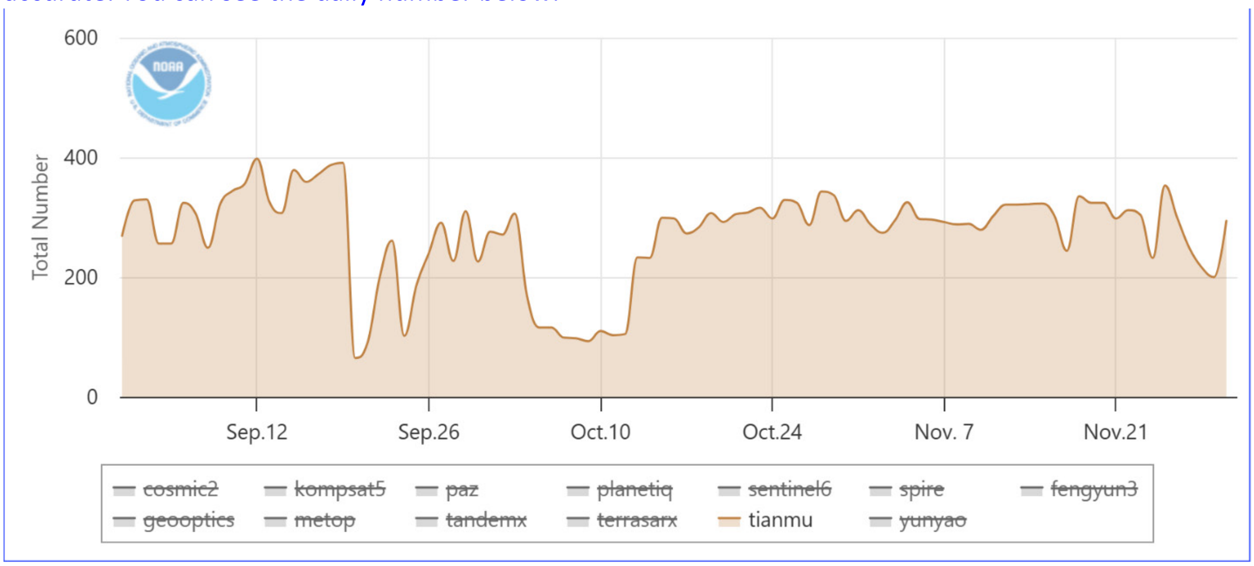

L213, etc but also elsewhere: The authors mention interpolation of certain quantities onto a common sampling grid. Since certain quantities such as position are non-trivial to interpolate to the required accuracy, please indicate summarily some details about the sufficiency of the interpolation. A quadratic is mentioned. Is a quadratic interpolation sufficient at 1Hz?

Reply:
We checked the code and actually the cubic spline interpolation instead of quadratic interpolation is used to align the satellite position vectors (GNSS and LEO), which are typically sampled at a low rate (e.g., 1 Hz), to the much higher excess-phase sampling rate (e.g., 50–100 Hz). For aligning 1Hz vectors to 50-100 Hz excess phase data, cubic spline interpolation (or high-order Lagrange polynomials) is the standard. The

interpolation is of the satellite ephemerides, which are inherently smooth. However, a quadratic fit may introduce jumps every second, which may create artificial noise in the high-rate phase residuals. A cubic spline ensures that not only is the position smooth, but the velocity and acceleration are also continuous at the boundaries. This mimics the physics of a satellite moving through a gravity field much more accurately. We have updated this in the text.

L234: "intersection of the line-of-sight with the WGS84 ellipsoid". The authors seem to be defining the SLTA=0 line and tangent point. This straight line between satellites is not a "line-of-sight", as the propagation trajectory is bent. Also, this line at this moment is not intersecting the ellipsoid, but is tangent to it. Please adjust the wording.

Reply:
We agree with this correction. The tangent point is defined as the location where the straight line between the transmitter and receiver touches the ellipsoid (straight line tangent altitude SLTA or height of straight line HSL equals 0). We adjusted the wording in Line 234 to be more physically accurate by replacing "intersection of the line-of-sight with the WGS84 ellipsoid" with "The tangent point location on the reference ellipsoid (WGS84) is determined where the straight line between the transmitter and receiver touches the ellipsoid"

L235: Presumably also the center of curvature has to be computed, and the coordinate system adjusted to be centered at that point.

Reply:
We agreed. We added "…and the center of curvature are computed at this reference tangent point…"

L257: Please provide some references to some of the existing studies around this truncation. This is a significantly complex issue, and the referencing should indicate its non-trivial character.

Reply:
We agree to add relevant references to existing studies on SNR-based signal truncation and cut-off methods to strengthen this part of the methodology: "An appropriate cut-off height is essential to ensure the accuracy and reliability of tropospheric information retrieved from GNSS signals in OL tracking mode (Sokolovskiy et al., 2009; Sokolovskiy et al., 2010; Adhikari et al., 2021; Paolella et al., 2025)."

L272-296: The MSIS-90 is ok for the temperature and to establish the approximate range in moist refractivity. However, one of the major issues is the presence of large vertical gradients of relative humidity, which may lead to large bending angles, superrefraction, etc. This model phase does not account for this possibility, nor seems to be considering the range of moistures. Is this model only identifying bit slips? Or will it introduce (false positive) bit corrections when there are RH gradients?

Reply:
The model phase derived from MSIS-90 (Lines 271–290) is an approximate, simplified reference for the sole purpose of identifying and correcting the abrupt, unphysical $\pm\pi$ navigation bit jumps (Lines 292–296). It is not intended to model high moisture gradients or super-refraction, and the final retrieval is not based on it.

L300: There are many parameters here. Although I agree that filtering is to some extent desirable, are these parameters optimal? Have other filtering parameters been tested, with the optimal results chosen? Under which criteria are the mentioned parameters optimal?

Reply:
The 0.5-second filter window is chosen because it approximately corresponds to the vertical extent of the first Fresnel zone (Lines 304–306), which is commonly used as a benchmark for the intrinsic vertical resolution of RO observations. The shorter windows (0.05 s (100 Hz) and 0.1 s (50 Hz)) are applied below 10 km to better resolve fine-scale structures in the lower troposphere while still maintaining sufficient noise suppression. The choice of these parameters is based on systematic tuning experiments in which multiple configurations were evaluated. The selected values showed here provide the most consistent statistical agreement with model outputs, such as ERA-5, yielding reduced bending angle bias and standard deviation compared with other tested configurations. We added these statements in the text.

L324: With MSP there may be a number of subcomponents that are stationary. There may be from none under anomalous propagation to several under multipath, although there is often just one dominant component. Should not L327 (eq 9) be a sum of perhaps several stationary components? Or has it been assumed that we focus only in the dominant? Does the text apply only to a case where there is one dominant component? Or does eq 9 express only one of the stationary components, in case there were several? Please clarify.

Reply:
Equation (9) is not a sum because it is presented within the context of applying the MSP approximation, which assume a single stationary point for a given frequency $\widehat{\omega}_q$ to illustrate the fundamental relationship between the phase derivative and the pseudo frequency. This assumption holds true for ideal occultations where the physics ensures a unique relationship between frequency and the ray's impact parameter. However, the FSI method is designed to overcome this limitation in realistic occultations (e.g., multipath). The technique uses the entire signal spectrum by applying a single Fourier transform to the entire complex signal, and applies corrections (such as resampling the signal with respect to open angle and correcting for radial variations) to eliminate interferences caused by higher-order derivatives and non-spherical trajectories. By applying these geometric and orbital corrections, the FSI method effectively modifies the signal so that the fundamental assumption of a single stationary point, or a unique frequency-to-ray relationship, is restored, allowing the single-component MSP approximation (and thus the framework containing Eq. 9) to be successfully applied to the corrected signal.

L363: Does this projection modify only the satellite orbits? Or, since the radius and center of curvature vary during the occultation, are these somehow reprojected also?

Reply:
The radius and center of curvature are treated as fixed parameters for a given RO event; they are not dynamically reprojected. The correction at Line 363 applies specifically to the satellite orbits, transforming the actual LEO/GNSS trajectories into idealized circular paths. This projection is performed relative to the pre-computed local center of curvature (Lines 364–370) to maintain a consistent coordinate system for the inversion process. We added these statements in the text.

L386: Does this criterion reject the low troposphere?

Reply:
The FSI amplitude criterion is the final step in a multi-stage SNR truncation process (Lines 240–261). Its purpose is to define the lowest valid bending angle-impact parameter pair for the profile. This threshold is indeed applied to reject signal contaminated by high noise (which often occurs in the lowest troposphere), as retaining this noisy data would lead to unphysical oscillations (Lines 257–260). This is an integral part of the noise-mitigation strategy. FSI method does provide an alternative estimation at all practical altitudes.

L399, eq 17: Aren't you using alpha in a different sense wrt before? Before it was the bending angle, while here it would suggest that alpha is the instantaneous bending angle minus the average bending angle (thus something like alpha-). Please make sure that the variable naming is consistent.

Reply:
We agreed. We revised the variable naming in the text and in Eq. (17) for consistency, to clearly define the variable representing the shifted bending angle relative to the central component.

L417: Please reference Eq 18, as it is not the definition of the neutral bending angle, but a good practical approximation to obtain it from alpha L1 and alpha L2 (Vorobev & Krasilnikova, 1994).

Reply:
We added a reference (Vorobev and Krasilnikova, 1994) for the dual-frequency linear combination (Eq. 18) and clarified its role as an approximation for the neutral atmosphere bending angle.

L441: Could this climatological data be specified?

Reply:
The bending angle profile is extended to 150 km using the MSIS-90 climatological model to ensure numerical stability during the Abel inversion's upper boundary condition.

L456…: Is this procedure valid only in hindcast? Or can it be applied in Near-Real-Time? (Availability of ERA5). It is also unclear in L457 if sigma_Year means the yearly variability of O-B, or the variability of B. In L466, thus later, it appears that it is O-B. Besides, is this statistic global? It could have been regional, by latitude band… Please clarify.

Reply:
The described QC is only for the post-processed RO data and cannot be applied to Near-Real-Time (NRT) processing, since ERA5 is not available in NRT. This QC is applied to the ROMEX dataset (hindcast) using ERA5 forecasts as the background (Lines 448–451). A separate internal QC system is under development for NRT processing (Lines 193–194, 474–475). σ_{year} is the annual mean standard deviation of the fractional bending angle O-B difference (%) (Lines 456–457). This statistic is based on the annual global average of the fractional bending angle O-B differences.

L470: Is this intended to analyze only height between 8-40km (or 10-40). The most interesting part is below 5 km. Is the method inappropriate here?

Reply:
The ERA5-based QC is intentionally restricted to the 10–40 km range (Lines 457–460) because this is where RO data quality and agreement with the background model are most stable. We do not apply this

specific QC below 10 km due to increased atmospheric variability and signal complexity (Lines 469–470). By bypassing this filter in the lower troposphere, we ensure that a larger volume of bending angle data is retained for the lower altitudes you mentioned.

L512: Comment here if the ionosphere was particularly active during the study period, thus if these are average circumstances, worse-case, etc.

Reply:
The ionospheric conditions during September-November 2022 were generally moderate and typical for the ascending phase of Solar Cycle 25, but intermittently disturbed (including several geomagnetic storms and enhanced irregularities). They are not representative of worst-case conditions, but also not purely quiet-average; they are better described as moderately active with episodic disturbances. We added these sentences in the text.

L514…: Comment here on the typical SNR of the different missions.

Reply:
Agreed. We added a table (Table 2) to summarize the typical SNR characteristics of GNSS receivers across different missions, supporting the discussion of mission-dependent performance and structural uncertainties.

Table 2 Signal-to-Noise Ratio (SNR)[#] characteristics of GNSS receivers across different missions

RO mission	GPS	GLONASS	Galileo	BeiDou
Metop B	730.2			
Metop C	789.3			
COSMIC-2	1295.1	1181.4		
SPIRE	387.9	707.2	316.4	
PlanetiQ	1440.5	1580.6	1124.8	1335.7
KOMPSAT-5	617.2			
PAZ	503.9			
TerraSAR-X	622.4			
TanDEM-X	549.3			

[#] Mean SNR between altitudes 60 km and 80 km, with unit (volt/volt)

L581: Isn't there a GO-to-FSI switch at a given altitude? 25km? It is mentioned again below, but would it not be appropriate to remind it somewhere in this section?

Reply:
Agreed. We explicitly stated in Section 3.2.2 that "Bending angles are computed using geometric optics (GO) above 25 km and the full spectrum inversion method below this altitude" (Line 311). We will point this out again in section 4.2.

L654…: The average and std between methods are shown, but not compared to the O-retrieval. Would it not be appropriate to indicate the fraction of the O-B differences that are explained by structural uncertainties of the retrieval? That is, are the O-B dominated by the retrieval method? Or is the exact retrieval a secondary source of uncertainty? This fraction is likely a function of altitude.

Reply:
We agree with the reviewer that explicitly framing the comparison between the total O-B uncertainty and the structural uncertainty is essential for the interpretation of our results. The comparison of the standard deviation profiles in Figs. 5-7 (total O-B Uncertainty) and Figs. 11-15 (structural uncertainty) is, in fact, the intended purpose of presenting the results in this manner, to isolate and quantify processing-related effects. The structural uncertainty is the portion of the total O-B error attributable solely to differences among retrieval algorithms.

To address the comment directly and clearly convey this key finding, we will add a sentence in the Discussions and Summary (Section 5) that explicitly: "The explicit comparison of the total O-B standard deviation (Figs. 5-7) with the structural uncertainty (Figs. 11-15) quantifies the contribution of retrieval algorithm differences to the total error budget, demonstrating that structural differences account for approximately one-fourth of the total O-B standard deviation over all the altitudes, providing a critical metric for interpreting ROMEX forecast impact studies and refining GNSS RO data assimilation systems."

L840: Does the FSI being presented restrict the vertical range where it is practical to apply? Some issues mentioned above would suggest that the procedure may be inappropriate for the low troposphere. Is it the case? Or the method does provide an alternative estimation at all practical altitudes?

Reply:
The FSI technique, a core component of the STAR FSI algorithm, is specifically designed to address the challenges posed by strong vertical gradients and multipath propagation that complicate retrievals in the lower troposphere. It is not inherently restricted in its vertical range and does provide an estimation at all practical altitudes. The reviewer's concern about the low troposphere is related to signal quality, not the FSI technique's applicability. The FSI method is explicitly chosen for its strength in resolving fine-scale atmospheric structures in the lower atmosphere. The SNR truncation criterion is a separate quality control step applied to the processed profile. Its purpose is to mitigate noise (which is often highest in the lowest troposphere) and prevent unphysical oscillations in the retrieved profiles. This quality control step ensures that only the lowest valid bending angle-impact parameter pair is retained, but the FSI method itself can produce estimates across the full vertical extent of the atmosphere.