https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Development and Evaluation of an AI Model for Simulating GNSS radio occultation Bending Angle through Ray Tracing

Xu Xu[1,2], Xiaolei Zou[3]

[1]Nanjing Innovation Institute for Atmospheric Sciences, Chinese Academy of Meteorological Sciences–Jiangsu Meteorological Service, Nanjing 210041, China
[2]Jiangsu Key Laboratory of Severe Storm Disaster Risk / Key Laboratory of Transportation Meteorology of CMA, Nanjing 210041, China
[3]Center of Data Assimilation for Research and Application, Nanjing University of Information Science & Technology, Nanjing, China

Correspondence to: Xiaolei Zou (xzou@nuist.edu.cn)

Abstract. Given computational cost and complexity of the ray-tracing observation operator and its adjoint operator for Global Navigation Satellite System (GNSS) Radio Occultation (RO) data assimilation, this study develops an AI model for simulating nonlocal bending angles using a Transformer encoder. The model uses ERA5 reanalysis data and COSMIC-2 RO geometry to learn the nonlocal mapping from atmospheric states along finite ray paths of bending angle simulated by ray tracing. A Transformer encoder extracts ray-path features, which are aggregated using target-dependent multi-head attention to predict bending angles. Training employs a composite loss function that combines weighted absolute errors below 8 km, layer-dependent relative error above 8 km, and constraints on layer-mean biases. The model is trained over East Asia from 30 June to 4 July 2024 and independently evaluated from 4 July to 8 July 2024. The AI model reproduces the vertical structure of bending angles simulated by ray tracing. The mean fractional difference is near 0% above 8 km, below 2% between 3 and 8 km, and within 8% below 3 km. Compared with the local Abel operator, the AI model exhibits less error dispersion. Incorporating tropopause height information reduces systematic bias near 18 km. Automatic differentiation also facilitates development of the AI model's adjoint. Relative sensitivity analyses of AI-simulated bending angles with respect to temperature and specific humidity further show that the AI model captures the nonlocal response patterns of the ray-tracing operator, with similar along-track distributions and sensitivity magnitudes. CPU timing results indicate that AI model and its adjoint calculations are approximately 1.5–3 times faster than ray tracing, demonstrating the potential for improved computational efficiency in forward ray-tracing and backward-adjoint calculations.

1 Introduction

Global Navigation Satellite System (GNSS) radio occultation (RO) provides high-vertical-resolution atmospheric observations with global coverage and all-weather capability. GNSS RO observations have become an important source of information for improving numerical weather prediction (NWP) through data assimilation. The RO observable, the bending angle, depends on the distributions of atmospheric refractivity and its gradients along the RO signal propagation path. Among the different

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

approaches for assimilating GNSS RO observations, direct assimilation of bending angles can be performed using a nonlocal ray-tracing observation operator, which provides a physically based representation of the measurement by explicitly calculating the propagation of the GNSS signal through a spatially varying atmospheric refractivity field. By minimizing the differences between simulated and observed bending angles, the atmospheric state, including the model refractivity field, can be adjusted through GNSS RO data assimilation.

Another widely used approach is to assimilate refractivity profiles retrieved from RO bending-angle observations through an Abel inversion under the assumption of local spherical symmetry. Refractivity is attractive for data assimilation because its observation operator is relatively simple and computationally inexpensive, and a number of studies have demonstrated positive impacts of refractivity assimilation on NWP analyses and forecasts. For example, using a local refractivity observation operator, Healy et al. (2005) assimilated CHAMP refractivity profiles into the Met Office NWP system and found improvements in upper-tropospheric and lower-stratospheric temperature analyses and forecasts, particularly in the Southern Hemisphere. However, because the assumption of local spherical symmetry neglects horizontal refractivity gradients along the RO ray path, retrieval errors can arise when strong horizontal refractivity gradients violate this assumption, particularly in the moist lower troposphere. Consequently, some of the nonlocal, spherically asymmetric refractivity information contained in the bending-angle observations may be lost when Abel-retrieved refractivity is assimilated.

An early effort to directly assimilate GNSS RO bending angles into NWP models was made by Zou et al. (1999), who developed a nonlocal ray-tracing observation operator together with its tangent-linear and adjoint operators for variational assimilation of GPS/MET bending-angle observations. The operator explicitly accounted for atmospheric refractivity and its spatial gradients along the ray path, thereby retaining the nonlocal characteristics of the RO measurement. Their adjoint sensitivity experiments further demonstrated that the influence of a bending-angle observation extends along the ray path rather than being confined to the tangent point. The ray-tracing operator was subsequently incorporated into the NCEP three-dimensional variational (3D-Var) analysis system (Zou et al., 2000) and applied to GPS/MET and CHAMP RO observations (Liu et al., 2001; Shao and Zou, 2002; Zou et al., 2004), demonstrating positive impacts on atmospheric analyses and forecasts. However, numerical ray tracing is considerably more complicated and computationally demanding than a local refractivity observation operator. Both 3D-Var and 4D-Var data assimilation systems require the development, maintenance, and repeated execution of the forward observation operator and its corresponding tangent-linear and adjoint operators when a conventional variational approach is used. Indeed, Zou et al. (2004) explicitly identified the development of a fast and operationally feasible method for RO assimilation as an outstanding challenge. These difficulties have posed challenges to the broader application of nonlocal ray-tracing bending-angle operators in operational RO data assimilation despite their physically attractive representation of the observations.

A progression from a computationally efficient Abel-transform observation operator toward a nonlocal ray-tracing operator occurred in the treatment of GNSS RO bending angles at ECMWF. Healy and Thépaut (2006) first implemented a local forward Abel-transform observation operator in the ECMWF four-dimensional variational (4D-Var) assimilation system. The forward Abel transform is computationally efficient but assumes local spherical symmetry of the model refractivity field.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Recognizing the limitations of this assumption, particularly in the lower troposphere, Healy et al. (2007) subsequently demonstrated the computational feasibility of the nonlocal ray-tracing observation operator using a fourth-order Runge–Kutta method to solve the differential equations governing the ray path through the refracting atmosphere. The nonlocal ray-tracing observation operator was later implemented for operational use at ECMWF (Healy, 2014).

Recent developments in artificial intelligence (AI), particularly deep learning (DL), have provided new approaches for constructing computationally efficient observation operators for data assimilation. For example, Lagerquist et al. (2021) used a DL emulator of the shortwave Rapid Radiative Transfer Model (RRTM) to achieve substantial computational acceleration while maintaining accuracy comparable to that of the original RRTM. Liang et al. (2022) developed a DL emulator of the Community Radiative Transfer Model (CRTM) for microwave satellite radiances and evaluated the accuracy of both simulated brightness temperatures and their Jacobians. Liang et al. (2023) further applied a DL radiance observation operator in an ensemble data assimilation system and demonstrated positive impacts of assimilating AMSU-A radiances. In addition to the computational efficiency of AI models relative to physically based observation operators, gradients of AI models can be obtained through automatic differentiation, potentially avoiding the complicated development and maintenance of tangent-linear and adjoint codes. These advances provide an opportunity to revisit the challenge identified by Zou et al. (2004) more than two decades ago.

Motivated by these developments, this study develops and evaluates an AI-based nonlocal GNSS RO bending-angle observation operator. The proposed AI model learns the nonlocal mapping from the atmospheric state to RO bending angles computed using the conventional ray-tracing operator. The accuracy and computational efficiency of the AI operator are evaluated against those of the conventional nonlocal ray-tracing operator. The gradients required for variational data assimilation are obtained from the AI operator through automatic differentiation and evaluated against those produced by the adjoint of the conventional ray-tracing operator. This paper focuses on the development and evaluation of the AI observation operator and its automatically derived gradients. A follow-on study will apply the AI model and its adjoint to GNSS RO data assimilation and assess its impacts on weather forecasts.

2 Data

2.1 COSMIC-2 RO Data and Study Region

This study uses the *atmPrf* atmospheric profile products provided by the COSMIC Data Analysis and Archive Center (CDAAC). These NetCDF products contain high-vertical-resolution (20 m) profiles of bending angle, impact height, refractivity, dry temperature and dry pressure, along with mean-sea-level geometric altitude, tangent-point coordinates, observation time, and RO geometry. Impact height, tangent-point latitude and longitude, geometric altitude, azimuth information, and radius of curvature from *atmPrf* are used to define the RO geometry for sampling atmospheric states from ERA5 and associating them with the corresponding bending-angle targets simulated by ray tracing. Impact height is used to

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

specify the bending-angle elevation level at the ray perigee and construct the target query, whereas geometric altitude is used to determine the actual sampling height within the ERA5 atmospheric column.

The data are organized into 6-h analysis intervals, with RO observations collected within 6-h time windows centered at each analysis time. Data from 0900 UTC on 30 June to 0900 UTC on 4 July 2024 (4 days) are used for training. Only profiles whose central locations fall within 9°–49°N and 86°–134°E are retained. An independent 4-day dataset covering 0900 UTC on 4 July to 0900 UTC on 8 July 2024 is used to evaluate the model's performance.

Figure 1 shows the topography of the study region and the locations and tangent directions of four representative RO profiles, with the tangent directions indicating the propagation directions of the RO signals. RO1, RO2, and RO3 are located over the South China Sea, near Taiwan, and near Hainan Island, respectively, whereas RO4 is located over mainland China. All four cases were observed on 30 June 2024, at 1419, 1420, 1424, and 1004 UTC, respectively. These cases represent different land-sea and topographic conditions and are used to examine the influence of horizontal atmospheric inhomogeneity on the simulated bending angles.

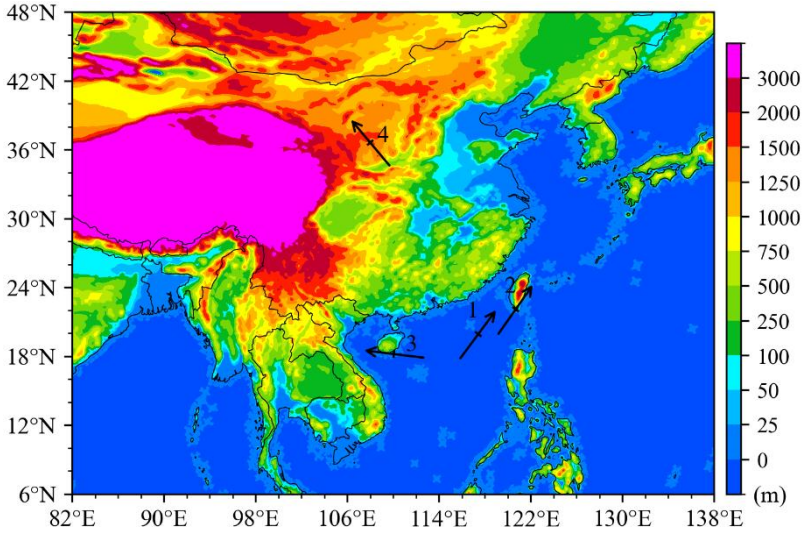


Figure 1. Terrain height (color shading) and four arbitrarily selected COSMIC-2 RO profiles, with the ray tangent directions indicated by 600-km-long black arrows and the perigee locations marked by small ticks. The observing times and perigee locations of the four RO profiles are as follows: 1419 UTC at (19.92°N, 117.43°E) for RO1, 1420 UTC at (22.09°N, 120.69°E) for RO2, 1424 UTC at (18.19°N, 110.06°E) for RO3, and 1004 UTC at (36.61°N, 108°E) for RO4, on 30 June 2024.

2.2 ERA5 Reanalysis

This study uses temperature and specific humidity from the 137 model levels of ERA5 reanalysis, together with surface geopotential and the natural logarithm of surface pressure. The vertical coordinate transformation follows the procedure described in the ECMWF Copernicus Knowledge Base for calculating pressure and geopotential on model levels. Pressure at each model level is reconstructed from the hybrid coordinate coefficients and surface pressure, while geopotential at each model level is calculated from the temperature and humidity profiles and surface geopotential. The geopotential is then

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

converted to geopotential height, which is subsequently converted to geometric height using the spherical Earth approximation. This resulting geometric height is used as the vertical coordinate for subsequent interpolation. Details are provided in the ECMWF technical document, "ERA5: compute pressure and geopotential on model levels, geopotential height and geometric height" (https://confluence.ecmwf.int/spaces/CKB/pages/158636068/ERA5+compute+pressure+and+geopotential+on+model+levels+geopotential+height+and+geometric+height).

Atmospheric refractivity is calculated from ERA5 temperature, specific humidity, and pressure as follows:

$$N = 77.6\frac{p}{T} + 3.73\times10^{5}\frac{e}{T^{2}} \quad (1)$$

where p and e denote atmospheric pressure and water vapor pressure, respectively. Water vapor pressure is calculated from specific humidity (q) as follows:

$$e = qp/(0.622 + 0.378q) \quad (2)$$

The first and second terms of Eq. (1) represent the contributions of dry air and water vapor, respectively. To collocate RO observations with ERA5 data, the nearest ERA5 grid point is selected horizontally, while linear interpolation is performed vertically with respect to geometric height.

2.3 Sample Construction

For each target impact height at 100-m intervals, the valid level in the *atmPrf* profile with the impact height closest to the target is selected. Using the tangent point and its geometric height as reference, 300 sampling points are uniformly distributed along the tangent direction within ±600 km of the tangent point. These points collectively define the sampling path. The geometric height of each sampling point is adjusted to account for Earth's local curvature, R , according to:

$$z(s) = \frac{R + z_0}{\cos(s/R)} - R \quad (3)$$

where z_0 denotes the geometric height at the central tangent point and s denotes the distance from the tangent point. ERA5 atmospheric states are then sampled at each point along the path. Each sampling path must contain at least 280 valid sampling points and missing atmospheric states along the path are estimated by linear interpolation.

The training dataset contains 633 RO profiles and 134,642 valid bending-angle samples. Each sample consists of 300 path points with associated atmospheric-state and geometric features, the target impact height, and the corresponding bending angle simulated by the ray-tracing operator. The ray-tracing-simulated bending angle serves as the training target for the AI operator to learn the nonlocal mapping represented by the ray-tracing operator.

The training samples are randomly divided to training (90%) and internal validation (10%) subsets. The independent test dataset covers the subsequent four-day period from 0900 UTC on 4 July to 0900 UTC on 8 July 2024. The internal validation set is used for model selection, while independent test set is used to evaluate the model's generalization performance.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

3 Methods

3.1 Ray-tracing Observation Operator

Zou et al. (1999) developed a ray-tracing observation operator, together with its tangent-linear and adjoint operators, for the direct assimilation of RO bending angles. The ray-tracing operator simulates the propagation paths of RO signals through a spatially inhomogeneous atmosphere by numerically integrating the ray equations. The two-dimensional ray equations can be written as:

$$\begin{cases} \dfrac{d^2\mathbf{r}}{d\tau^2} = n\nabla n \\ d\tau = \dfrac{ds}{n} \end{cases} \tag{4}$$

Where the variables r, s, n, and τ denote the ray position vector, arc length along the ray, atmospheric refractive index, and ray parameter, respectively. By introducing the ray tangent vector $\mathbf{y} = d\mathbf{r}/d\tau$, the second-order equation can be rewritten as a system of first-order equations:

$$\begin{cases} \dfrac{d\mathbf{r}}{d\tau} = \mathbf{y} \\ \dfrac{d\mathbf{y}}{d\tau} = n\nabla n \end{cases} \tag{5}$$

Numerical integration of this system yields the ray trajectory and propagation direction. The bending angle is then calculated from the ray tangent vectors at the GNSS and low-Earth-orbit (LEO) satellite positions:

$$\alpha_{\text{Ray}} = \angle(\mathbf{t}_G, \mathbf{t}_L) \tag{6}$$

Where the two vectors $\mathbf{t}_G$ and $\mathbf{t}_L$ denote the ray tangent vectors at the GNSS and LEO satellite positions, respectively.

Because the ray trajectory and its directional changes depend on both refractivity and its spatial gradients along the propagation path, ray-tracing bending angles exhibit nonlocal characteristics. Their responses to atmospheric temperature and humidity variations can extend along the entire propagation path. In this study, ray-tracing bending angles calculated using the ERA5 reanalysis data serve as training targets for the AI model, which learns to approximate the mapping between atmospheric states and nonlocal bending angles.

The adjoint of the ray-tracing operator is constructed from the tangent linear formulation of the forward operator (Zou, 2020). By propagating sensitivities backward through the computational process, the adjoint efficiently calculates the gradients of simulated bending angles with respect to temperature and specific humidity (Zou et al., 1999). These gradients are used for evaluating relative sensitivities and computing the gradients of cost functions for minimization in variational data assimilation.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

3.2 Abel Transform Operator

The Abel transform is based on the assumption of local spherical symmetry, under which the atmospheric refractive index depends only on the radial distance r from Earth's center of curvature. Under this assumption, the impact parameter a remains constant along the ray, i.e., $a = nr\sin\theta$, where θ is the angle between the ray tangent and the radial direction. Defining the refractive radius as $x = nr$ and assuming that it increases monotonically with r, the bending angle can be calculated from the refractive index using the forward Abel transform (Fjeldbo et al., 1971):

$$\alpha_{\text{Abel}}(a) = -2a\int_a^{\infty} \frac{d\ln n / dx}{\sqrt{x^2 - a^2}} dx \tag{7}$$

Unlike the ray-tracing operator, which explicitly accounts for atmospheric inhomogeneity along the propagation path, the Abel operator approximates the atmosphere as locally spherically symmetric, thereby neglecting horizontal refractivity gradients. Consequently, its bending-angle estimates may exhibit substantial errors in regions of strong horizontal atmospheric inhomogeneity. Furthermore, the conventional Abel transform requires the refractive radius to increase monotonically with radial distance. In this study, bending angles calculated using the forward Abel transform serve as an additional reference for evaluating the AI model against the ray-tracing operator.

3.3 AI Model for Ray Tracing

This study develops a Transformer-based AI surrogate model for the nonlocal bending-angle observation operator. Because bending angles depend on the cumulative effects of atmospheric refractivity along the ray path, the model must capture the contributions of atmospheric conditions at different locations along the path. The Transformer architecture, with its self-attention mechanism, can capture long-range dependencies among sampling points along the ray path (Vaswani et al., 2017), making it suitable for learning the nonlocal mapping between atmospheric conditions and bending angles. The model encodes atmospheric and geometric features sampled along the tangent path using a Transformer encoder and predicts bending angles through target-dependent multi-head attention-based path aggregation. Bending angles simulated by ray tracing serve as training targets, and the model is evaluated using data from an independent time period. The following subsections describe the input features, network architecture, training objective, and experimental configuration.

3.3.1 Input Features and Tropopause Information

A total of 18 input features are used in this study (Table 1), including atmospheric state variables, refractivity gradients and their absolute values, path geometry, vertical height, and tropopause information. Atmospheric state features characterize the conditions that determine atmospheric refractivity; gradient features describe the local variations responsible for ray bending and their distribution along the propagation path; path geometry features represent sampling locations, propagation directions, and integration scales; and vertical height and tropopause information provide additional information on atmospheric vertical

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

structure. Together, these features enable the model to learn the nonlocal mapping between atmospheric states along the sampling path and the target bending angle.

To provide the model with information on the sharp changes in atmospheric thermal structure near the tropopause, tropopause height is included as an additional input feature. The tropopause height is diagnosed from the temperature lapse rate at the nearest ERA5 grid point to the RO tangent point. The algorithm searches upward from 5.5 km to identify the lowest candidate level at which the local lapse rate does not exceed 2 K km^{-1}. It then verifies that the mean lapse rate over approximately the next 2 km above that level also satisfies this criterion, allowing a tolerance of 0.4 km to account for differences between the discrete model-level spacing and the 2-km verification interval. This procedure approximates the WMO tropopause criterion using discrete model levels.

Table 1. Transformer model input.

No.	Variable	No.	Parameter	No.	Parameter
1	T	7	$\angle(\nabla N, \mathbf{t}_{perigee})$	13	Δx
2	q	8	s	14	Δy
3	N	9	Δs	15	Δz
4	$\partial N/\partial z$	10	t_x	16	$h_{a,\mathrm{local}}$
5	dN/ds	11	t_y	17	z
6	$\partial^2 N/\partial z^2$	12	t_z	18	$z_{tropopause}$

3.3.2 Architecture of the Transformer-Based AI Observation Operator

The model follows the path-integration principle of the ray-tracing operator and uses a Transformer to learn the nonlocal relationships between atmospheric conditions sampling points along the path and the target bending angle. For each target impact height (h_a), a sequence of path features is constructed from 300 sampling points along the tangent path. A Transformer encoder is then used to capture nonlocal dependencies among the sampling points. The input features at each sampling point are first projected into a 128-dimensional latent space, to which sinusoidal positional encodings are added. The encoder consists of three Transformer layers, each with four self-attention heads and a feed-forward network with a hidden dimension of 256. The positional encodings and explicit path coordinates jointly represent the sampling order and spatial locations.

To establish the relationship between the target bending angle and the path features, the target impact height is encoded as a seven-dimensional vector comprising five height-dependent features (normalized height, exponential decay, inverse distance, logarithmic height, and a low-level Gaussian weight) and two tropopause-relative features (a signed distance and a distance-based weight). This target encoding is then used to generate the query vectors. The path encodings generate key vectors and

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

learn path contributions for each integration head. The model employs four integration heads, each of which computes attention weights and aggregates path contributions. The outputs of the four heads are then combined using target-dependent fusion coefficients to obtain the predicted bending angle:

$$\hat{\alpha}(h_a) = \sum_{m=1}^{M} \rho_m(h_a) \sum_{i=1}^{L} w_{mi}(h_a) \varphi_{mi}(\mathbf{x}_i, h_a) \tag{8}$$

where M=4 is the number of integration heads, L=300 is the number of path sampling points, x_i denotes the input features at the i-th path point, w_{mi} is the path attention weight for the m-th integration head, ρ_m is the corresponding head fusion coefficient, and represents the learned path contribution.

The attention weight for the m-th integration head is defined as follows:

$$w_{mi} = softmax_i\left(\frac{\mathbf{q}_m \cdot \mathbf{k}_{mi}}{\sqrt{64}} + \ln \Delta l_i\right) \tag{9}$$

where $\mathbf{q}_m$ and $\mathbf{k}_{mi}$ denote the query and key vectors, respectively, each with a dimension of 64, and Δl_i is the path integration step length. By incorporating the sampling interval into the attention-weight calculation, the model accounts for both learned positional relevance and path-integration scale during path aggregation. Each integration head independently aggregates the path contributions, and the resulting outputs are combined using target-dependent fusion coefficients to obtain the normalized bending-angle prediction, which is then transformed back to radians through inverse normalization.

The adjoint of the AI observation operator is implemented using PyTorch's automatic differentiation framework. The derivatives of the predicted bending angles with respect to temperature and specific humidity are computed by differentiating the entire computational chain, from atmospheric feature construction to Transformer prediction. This allows the AI operator to provide both forward bending-angle simulations and the corresponding gradient information required for relative sensitivity analysis and variational data assimilation.

3.3.3 Loss Function and Training Configuration

To account for the pronounced decrease in bending-angle magnitude with increasing height, a composite loss function is adopted, consisting of a low-level weighted squared-error term, a piecewise relative-error term for the middle and upper atmospheric levels, and a layer-averaged bias constraint. This formulation aims to balance simulation accuracy across different heights while suppressing systematic bias.

The low-level weighted squared-error term L_{low} is based on the squared bending-angle error ($(\alpha_{AI} - \alpha_{Ray})^2$), with a height-dependent weighting function to emphasize errors below 8 km. The piecewise relative-error term L_{upper} is based on the relative bias ($(\alpha_{AI} - \alpha_{Ray}) / \alpha_{Ray}$). The mean loss is calculated separately for the 8–14, 14–20, and 20–30 km intervals, with weights of 0.40, 0.40, and 0.20, respectively, to constrain relative simulation accuracy across these height ranges. The layer-averaged bias constraint L_{bias} calculates the mean relative bias within four height intervals: 0–8, 8–20, 20–30, and above 30 km. The squared

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

mean biases are then weighted by 0.02, 0.02, 0.025 and 0.025, respectively, to penalized systematic overestimation or underestimation. The total loss function is defined as:

$$L = 0.42L_{low} + 1.16L_{upper} + L_{bias} \tag{10}$$

The model is trained using the AdamW optimizer with a cosine annealing learning-rate scheduler. The initial learning rate is 5×10^{-5}, the weight decay coefficient is 10^{-4}, the gradient-norm clipping threshold is 1, and the dropout rate is 0.1. Training is performed for 45 epochs using distributed data parallelism across two nodes and four processes, with a batch size of 32 per process and a global batch size of 128. The model with the lowest validation loss is selected for evaluation on the independent test dataset.

3.3.4 Experimental Design for Assessing Values of Tropopause Information

To investigate the influence of tropopause information on bending-angle simulations, the AI model incorporating tropopause information (AI) is compared with a control model that excludes this information (AI_0). The AI model uses the relative positions of path points and the target impact height with respect to the tropopause to represent the atmospheric stratification, whereas AI_0 excludes these inputs. Both models are evaluated against two-dimensional ray-tracing results using the same test dataset. The comparison examines individual bending-angle profiles, the mean fractional difference and its standard deviation at each height, with particular emphasis on systematic biases near and above the tropopause and changes in low-level errors.

3.4 Relative Sensitivity to Temperature and Specific Humidity

To compare the sensitivities of simulated bending angles by the ray-tracing and AI operators to temperature and specific humidity, rays from RO1 with geometric tangent heights near 15, 10, 5, and 2.5 km are selected. Taking the bending angle as the response variable, the dimensionless relative sensitivities at each background grid point are defined as follows (Zou, 2020):

$$\mathbf{s} = \frac{\mathbf{x} \odot \nabla_{\mathbf{x}} \alpha}{\alpha} \tag{11}$$

where $\mathbf{x} = [\mathbf{T}, \mathbf{q}]^T$ is the atmospheric state vector, $\nabla_{\mathbf{x}} \alpha$ is the gradient of α with respect to $\mathbf{x}$, as computed using the adjoint operator, and $\odot$ denotes the element-wise product. Relative sensitivity quantifies the fractional change in bending angle in response to a fractional perturbation in temperature or specific humidity. Its sign indicates the direction of the bending-angle response to the corresponding perturbation.

4 Results

This section evaluates the AI model's ability to reproduce ray-tracing bending angles and the sensitivities of these bending angles to temperature and specific humidity. The evaluation includes comparisons of bending angles from the AI model, the ray-tracing operator, and the Abel operator for individual RO cases, along with analyses of refractivity distributions and ray

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

trajectories to explain the differences; an examination of tropopause height variability and its implications for AI model performance; height-dependent error statistics from independent test samples, with emphasis on the impact of tropopause information; comparisons of relative sensitivities to temperature and specific humidity between the AI model and the ray-tracing operator; and an assessment of computational efficiency and limitations for practical data assimilation.

4.1 Case Analysis

Figure 2 compares the bending angles simulated by the ray-tracing operator (α_{Ray}), the Abel operator (α_{Abel}), and the AI models with (α_{AI}) and without (α_{AI_0}) tropopause information for four RO cases (RO1–RO4). The bending angles simulated by all four methods decrease with increasing impact height and are nearly identical above 10 km. At lower heights, the bending angles increase rapidly with decreasing impact height and exhibit pronounced local vertical variations. Distinct variations in bending angle occur near 3 km for RO1, 4–5 km for RO2, 4–6 km for RO3, and 5–6 km for RO4, reflecting differences in the lower-atmospheric structures among the cases. Both AI models generally reproduce these vertical variations.

In terms of fractional differences, the Abel and AI operators yield bending angles close to the ray-tracing results at mid- and upper-tropospheric levels, with the main discrepancies concentrated in the lower troposphere, where bending angles vary more sharply. For RO1, the Abel operator shows a substantial negative bias near 3 km. The AI simulations are closer to the ray-tracing results for RO2 near 5 km and RO4 over approximately 6–9 km, whereas the errors of the different methods alternate in sign over 4–7 km for RO3. Above 15 km, the overall differences are small, with the AI model exhibiting the smallest differences. For AI_0, the relative differences in RO1 and RO2 increase substantially near 17.5 km, while pronounced biases occur above 18 km for RO3 and RO4. These discrepancies may be associated with the absence of explicit tropopause information, which could limit the model's ability to capture the effects of sharp atmospheric gradients near the tropopause on predicted bending angles. Overall, the AI model generally provides more accurate bending-angle simulations than the Abel operator across the four RO cases, although the magnitude of improvement depends on height and atmospheric conditions. The remaining low-level errors may reflect the difficulty of representing sharp, localized atmospheric variations.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

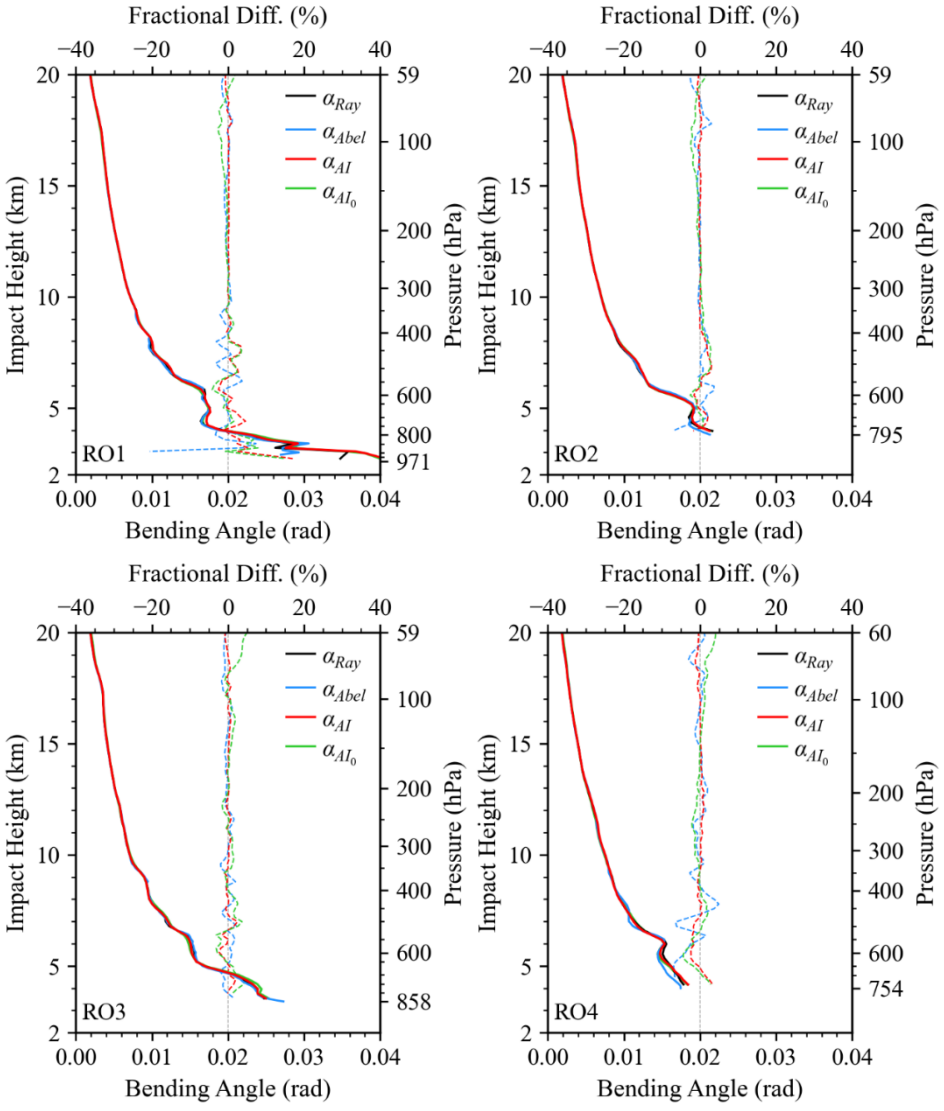


Figure 2. Vertical profiles of RO bending angle for ROs 1–4 obtained from raytracing simulation (α_{Ray} , solid black curve), the Abel transform (α_{Abel} , solid blue curve), AI-model simulations with tropopause height as an input (α_{AI} , solid red curve), and AI-model simulations without tropopause height as an input (α_{AI_0} , solid green curve), together with the corresponding fractional differences: $(\alpha_{Abel}-\alpha_{Ray})/\alpha_{Ray}$ (dashed blue curve), $(\alpha_{AI}-\alpha_{Ray})/\alpha_{Ray}$ (dashed red curve), and $(\alpha_{AI_0}-\alpha_{Ray})/\alpha_{Ray}$ (dashed green curve).

To investigate the physical factors underlying these differences, Figure 3 shows the horizontal distribution of refractivity at 850 hPa and vertical cross sections of refractivity along the RO tangent directions for RO1 and RO2 at 1200 UTC on 30 June 2024. Refractivity exhibits pronounced horizontal inhomogeneity across the study region. This is particularly evident at lower levels, where the water-vapor contribution to refractivity varies substantially in both the horizontal and vertical directions, whereas the dry-air contribution is comparatively smooth. The refractivity contours in the vertical cross sections show

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

pronounced undulations and asymmetric patterns, indicating that the atmospheric conditions encountered along the ray path cannot be fully represented by a single vertical profile near the tangent point. Such inhomogeneity along the ray path is an important physical source of differences between bending angles simulated by the Abel and ray-tracing operators, owing to the spherical-symmetry assumption underlying the Abel operator.

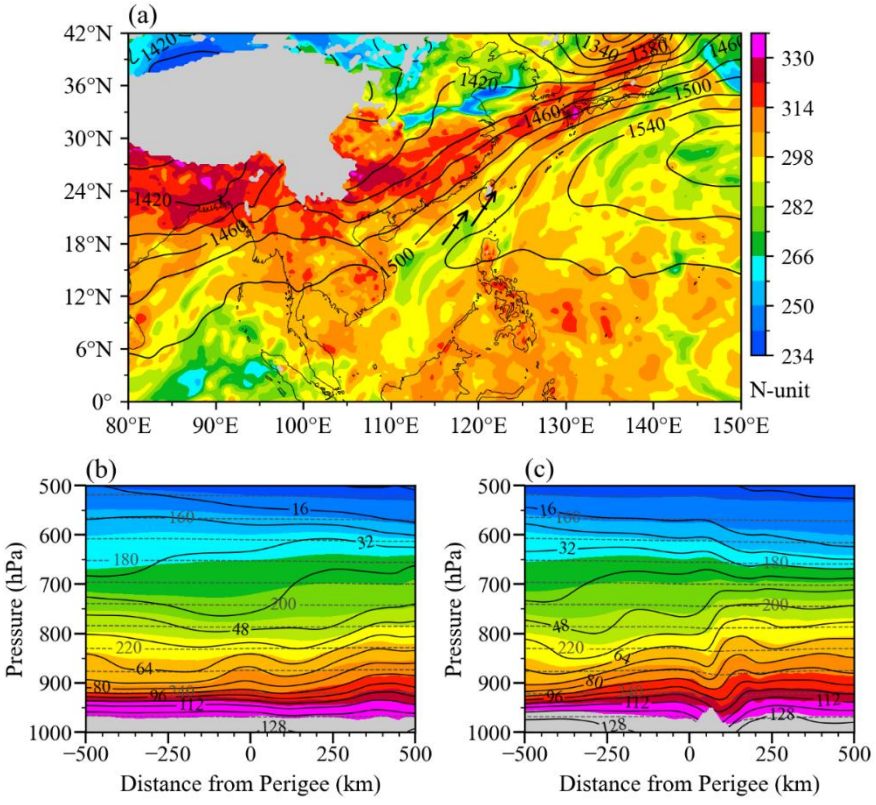
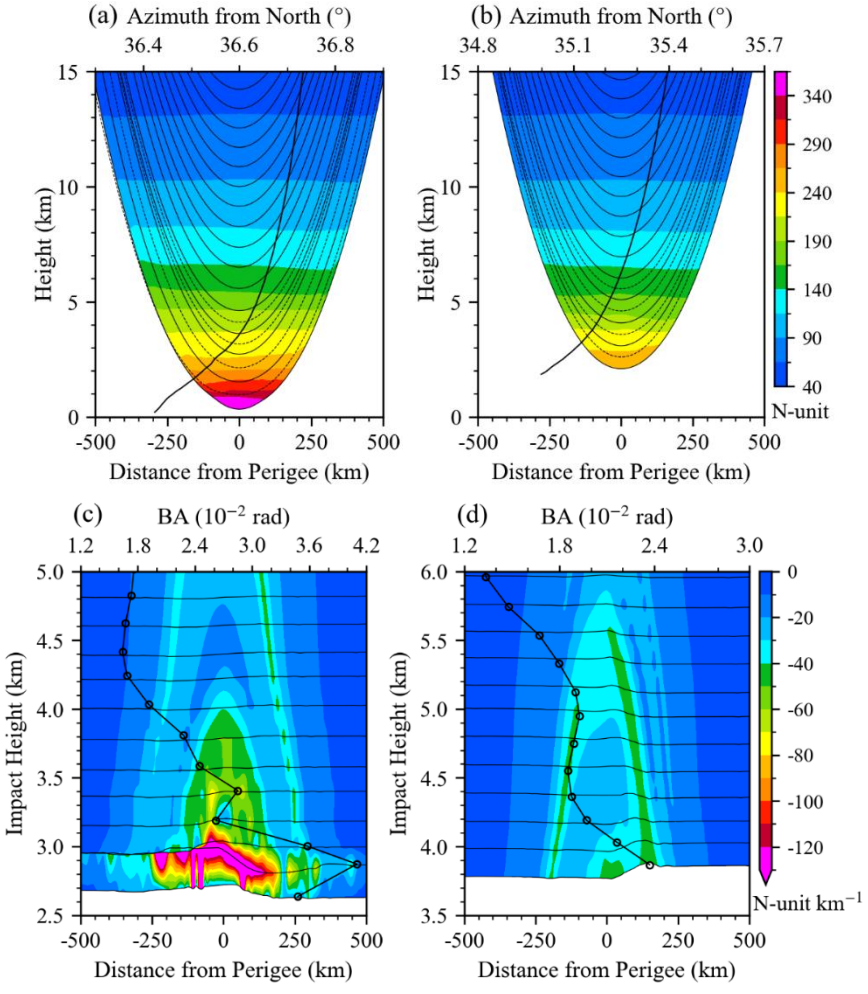


Figure 3. (a) Horizontal distribution of refractivity (color shading) and geopotential height (black contours, units: gpm) at 850 hPa, and (b)–(c) vertical cross sections of refractivity (color shading), its dry-air contribution (N_{dry} , dashed black curves), and water vapor contribution (N_{vapor} , solid black contours) along the ray paths of (b) RO1 and (c) RO2, calculated from the ERA5 reanalysis at 1200 UTC on 30 June 2024.

Figure 4 further illustrates the ray paths and the corresponding refractivity-gradient structures for RO1 and RO2. The rays reach their minimum heights near the tangent points and gradually ascend toward both ends. Consequently, rays with different impact heights are influenced by refractivity structures across different spatial regions. In both cases, the vertical refractivity gradient exhibits localized enhancements and pronounced horizontal variations at lower levels. These variations are particularly prominent near 3 km for RO1, corresponding to the distinct bending-angle variation at approximately the same height in Fig. 2. In comparison, RO2 exhibits smaller magnitudes of vertical refractivity gradients and weaker spatial variations, corresponding to a smoother bending-angle profile. The horizontal refractivity gradients are approximately two orders of

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

magnitude smaller than the vertical gradients, based on their respective magnitudes, and show relatively little difference between RO1 and RO2.



335

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

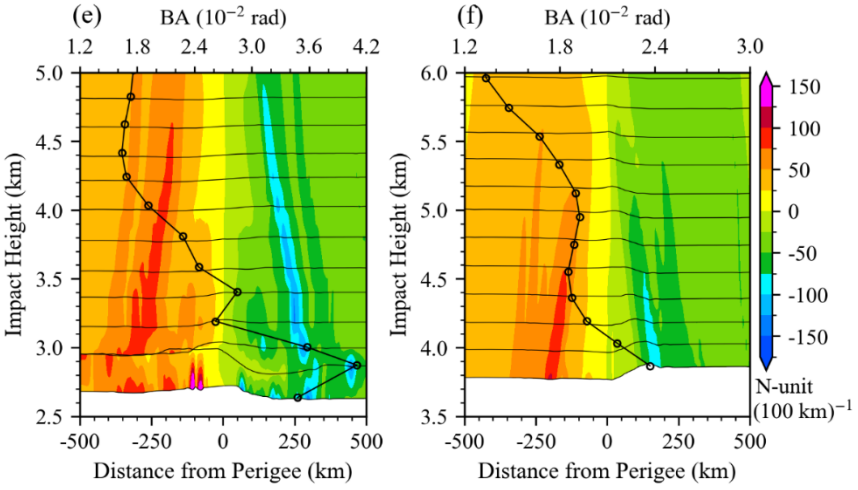


Figure 4. (a)–(b) Simulated ray paths (thin black solid curve at 800-m vertical interval at perigee, thin dashed black curve in the middle of two neighbouring solid curves), the azimuth angle of the 2D occultation plane with respect to the north (thick black curve), and refractivity (color shading); (c)–(d) variations of simulated ray paths with respect to impact height (thin black curve), raytracing-simulated bending angle (open circles), and vertical gradients of refractivity (N-unit km^{-1}, color shading), and (e)–(f) same as (c)–(d) except for horizontal gradients of refractivity along the rays (N-unit (100 km)$^{-1}$, color shading), within ±500-km horizontal distances from the perigee for RO1 (left panels) and RO2 (right panels).

These results indicate that local variations in low-level bending angles are closely associated with the complex refractivity-gradient structures encountered along the ray paths. Compared with the Abel operator, which assumes spherical symmetry and uses a single vertical refractivity profile, the AI model incorporates horizontal atmospheric inhomogeneity through finite-path sampling of atmospheric states and refractivity gradients, enabling it to more accurately reproduce the ray-tracing results in some cases.

4.2 Tropopause Height Variability and Its Relevance to the AI Model

Figure 5 shows the temporal distribution of tropopause heights for the 633 RO profiles used in training. Most tropopause heights are concentrated between 15 and 18 km, with median values generally ranging from 16 to 17 km. However, the interquartile ranges and outliers vary across analysis times, and the tropopause heights of some profiles are as low as approximately 11–12 km. This indicates that even within the relatively short training period, the tropopause heights associated with different RO profiles exhibit substantial variability. Therefore, a fixed height cannot fully characterize the vertical position of the tropopause across different profiles.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

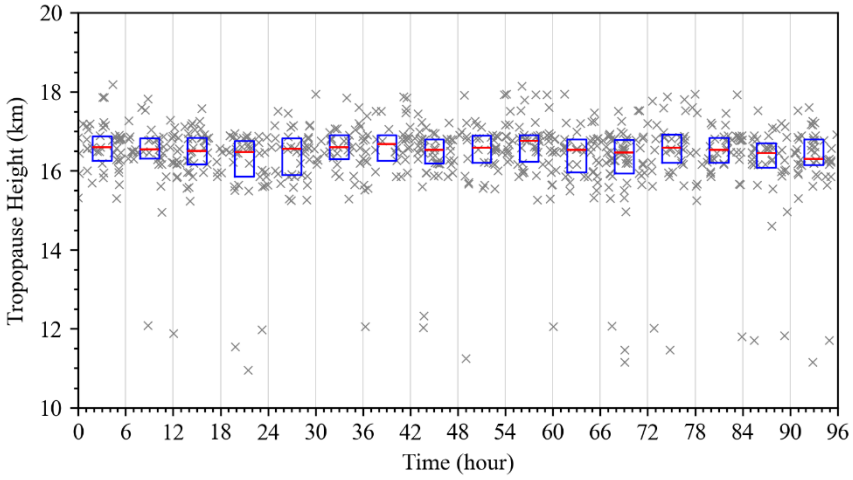


Figure 5. Scatter plot of tropopause heights for 633 RO profiles available during the AI training period from 0900 UTC on 30 June to 0900 UTC on 4 July 2024 over the region bounded by 9°–39°N, 86°–134°E. The blue boxes represent the interquartile ranges (25th–75th percentiles), and the red lines indicate the median tropopause heights of RO profiles observed within ±3 h of each analysis time, with analysis times separated by thin gray vertical lines.

Figure 6 further illustrates the spatial distribution of the training samples, the RO tangent directions, and the corresponding tropopause heights. The RO observations cover China and the surrounding seas, with diverse tangent directions and spatial locations. Tropopause heights exhibit pronounced spatial variations across the region, with relatively low values for some profiles in the northern and northeastern regions and higher values for some profiles in the west. Together with the 500-hPa geopotential height field, these distributions suggest that the training samples encompass diverse synoptic-scale circulation patterns and atmospheric stratification conditions. Because the 500-hPa circulation field represents only a single time, whereas the RO observations span the entire training period, the figure is intended primarily to illustrate the spatial representativeness of the samples rather than establish a quantitative relationship between tropopause height and instantaneous circulation. Taken together, Figures 5 and 6 illustrate the variability in tropopause height across the training samples, suggesting that tropopause-relative height could provide the AI model with information about vertical atmospheric stratification beyond geometric height, potentially helping it distinguish bending-angle variations under different background conditions.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

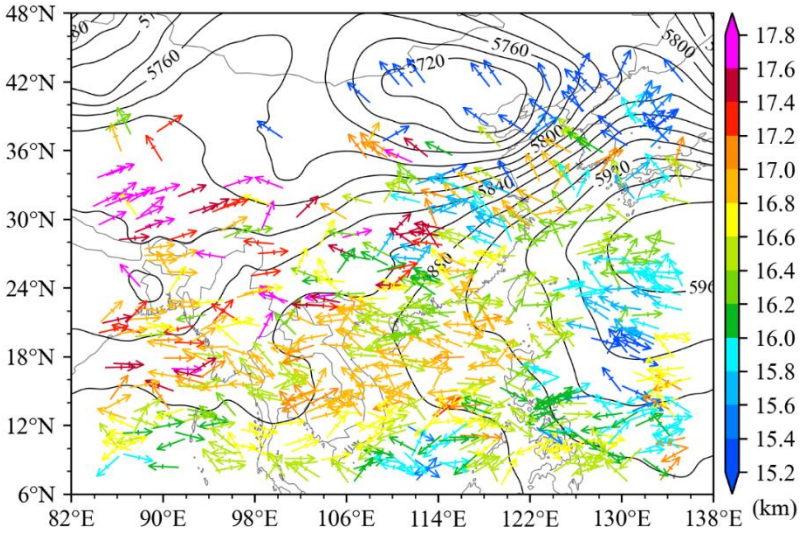


Figure 6. Spatial distribution of ray tangent directions (600-km-long black arrows) and perigee locations (small ticks at the midpoints of the arrows) of RO profiles during the training period, with their tropopause heights indicated by colors. Also shown is the ERA5 500-hPa geopotential height field (black contours, units: gpm) at 1200 UTC on 2 July 2024.

4.3 Statistical Evaluation

To evaluate the simulation accuracy of the AI model, an independent testing is conducted using RO observations from 0900 UTC on 4 July to 0900 UTC on 8 July 2024. Figure 7 shows the bending-angle profiles simulated by AI, AI_0, and the Abel operator as functions of impact height, together with their fractional differences relative to the ray-tracing results. The left panels show that all methods reproduce the overall decrease in bending angle with increasing impact height. The spread among individual profiles increases markedly toward lower heights, reflecting substantial variations in lower-atmospheric conditions and refractivity structures. In contrast, the fractional differences in the right panels more clearly reveal the simulation errors of the different operators. For AI, the fractional differences at mid- and upper levels are concentrated near zero, with the narrowest overall distribution. Height-dependent biases are evident at approximately 5–8 km, while the error distribution broadens considerably at the lowest levels (approximately 2.5–3 km), where some profiles exhibit large positive fractional differences.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

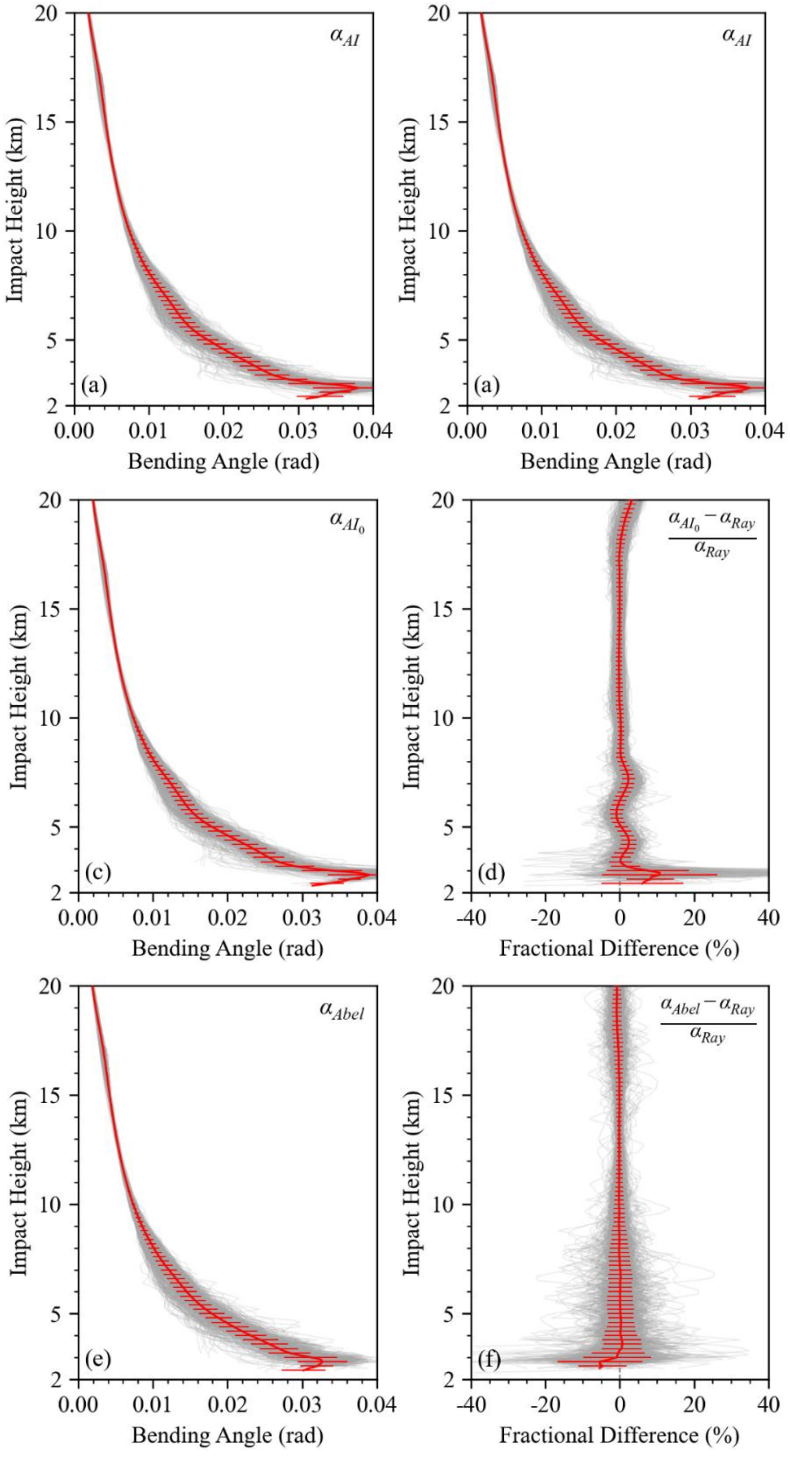


Figure 7. Spaghetti maps of vertical profiles of α_{AI}, α_{AI_0} and α_{Abel} (left panels) and their fractional differences $(\alpha_{AI}-\alpha_{Ray})/\alpha_{Ray}$, $(\alpha_{AI_0}-\alpha_{Ray})/\alpha_{Ray}$, and $(\alpha_{Abel}-\alpha_{Ray})/\alpha_{Ray}$ (right panels), for RO profiles during the testing period from 0900 UTC on 4 July to 0900 UTC

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

on 8 July 2024. Also shown are the mean (red curve) and standard deviation ($\pm\sigma$, red horizontal bars) calculated within each 0.2-km impact-height interval.

Compared with AI, AI$_0$ exhibits a more pronounced positive bias around 17–20 km, indicating a systematic overestimation of bending angles at upper levels when tropopause information is excluded. The Abel operator generally exhibits small fractional differences at mid- and upper levels, but its fractional differences are more widely scattered at lower levels, with some profiles exhibiting substantial positive or negative deviations. These results indicate that the AI model maintains high simulation accuracy at mid- and upper levels while reducing errors for some individual profiles at lower levels. The influence of tropopause information is primarily reflected in the reduction of upper-level biases.

Figure 8 further summarizes the means and standard deviations of the bending angles and their fractional differences. In Fig. 8a, the mean bending-angle profiles produced by all operators largely overlap. The standard deviation increases toward lower heights, with a marked increase near the lowest levels. This indicates that all operators reproduce the main vertical variations in bending angle and the variability among samples. Figure 8b shows that the mean fractional difference for AI is close to zero above approximately 8 km, whereas AI$_0$ exhibits a more pronounced positive bias at upper levels. Both AI and AI$_0$ exhibit height-dependent biases around 5–8 km, and their mean fractional differences become distinctly positive at approximately 2.5–3 km, indicating systematic overestimation at these heights. The Abel operator has small mean biases between 2.5 and 8 km but exhibits a negative bias at the lowest levels. The standard deviation of the fractional differences generally increases toward lower heights. Below 10 km, the Abel operator exhibits greater dispersion than the AI model, while the standard deviations of all methods increase markedly below 2 km.

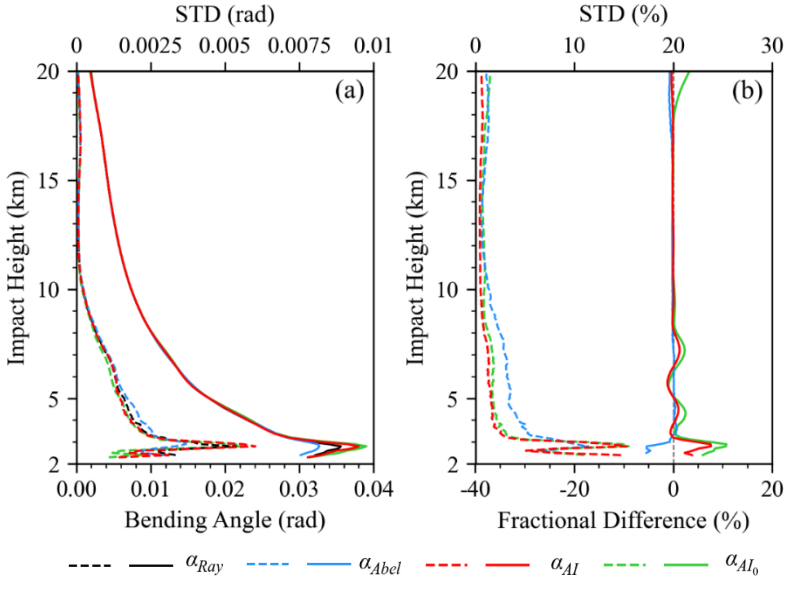


Figure 8. (a) Mean (solid curves) and standard deviation (dashed curves) of the vertical bending-angle profiles dobtained from raytracing simulations (α_{Ray} , black), the Abel transform (α_{Abel} , blue), and AI models with (α_{AI} , red) and without (α_{AI_0} , green) tropopause height as an input. (b) Mean (solid curves) and standard deviation (dashed curves) of vertical profiles of the fractional differences, $(\alpha_{AI} - \alpha_{Ray}) / \alpha_{Ray}$,

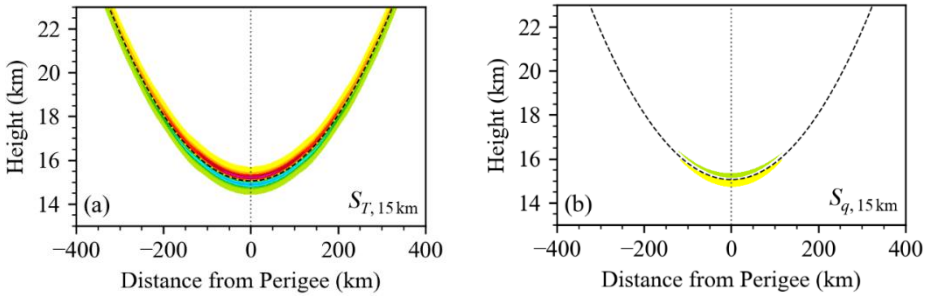
https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

$(\alpha_{AI_0} - \alpha_{Ray}) / \alpha_{Ray}$, and $(\alpha_{Abel} - \alpha_{Ray}) / \alpha_{Ray}$, for RO profiles during the testing period from 0900 UTC on 4 July to 0900 UTC on 8 July 2024.

4.4 Relative Sensitivity Analysis for Temperature and Specific Humidity

In addition to bending-angle simulation accuracy, the sensitivity of the AI model to the background atmospheric state is an important criterion for evaluating its potential to replace the ray-tracing operator. For RO1, Figures 9 and 11 show the relative sensitivities of the bending angles to temperature and specific humidity computed using the ray-tracing operator and the AI model, respectively, at geometric heights of approximately 15, 10, 5, and 2.5 km. To reduce local numerical fluctuations arising from grid discretization, two-dimensional Gaussian smoothing is applied to the relative sensitivity fields for temperature and specific humidity (Babaud et al., 1986), with horizontal and vertical smoothing scales of 12 and 0.15 km, respectively. Valid-data normalization is used during smoothing to prevent missing values from affecting neighbouring valid grid points. The positive and negative sensitivity amplitudes are then adjusted separately to preserve the original positive and negative extrema. This smoothing is used solely for visualization and does not alter the original relative sensitivity calculations.

Figure 9 shows that the relative sensitivities of bending angles to temperature and specific humidity sensitivities computed using the ray-tracing operator primarily form continuous response bands along the ray paths, reflecting the nonlocal dependence of bending angles on atmospheric conditions along the propagation paths. The temperature sensitivity exhibits a pattern of positive responses above and negative responses below the ray path, with relatively continuous spatial distributions at all target heights. In contrast, the specific-humidity sensitivity exhibits a stronger dependence on target height. At approximately 15 km, it is relatively weak and concentrated mainly near the tangent point. As the target height decreases, the sensitivity strengthens and extends farther along the path in both directions. The spatial distributions are particularly pronounced near 5 and 2.5 km, highlighting the important influence of lower-tropospheric water vapor on bending angles.



https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

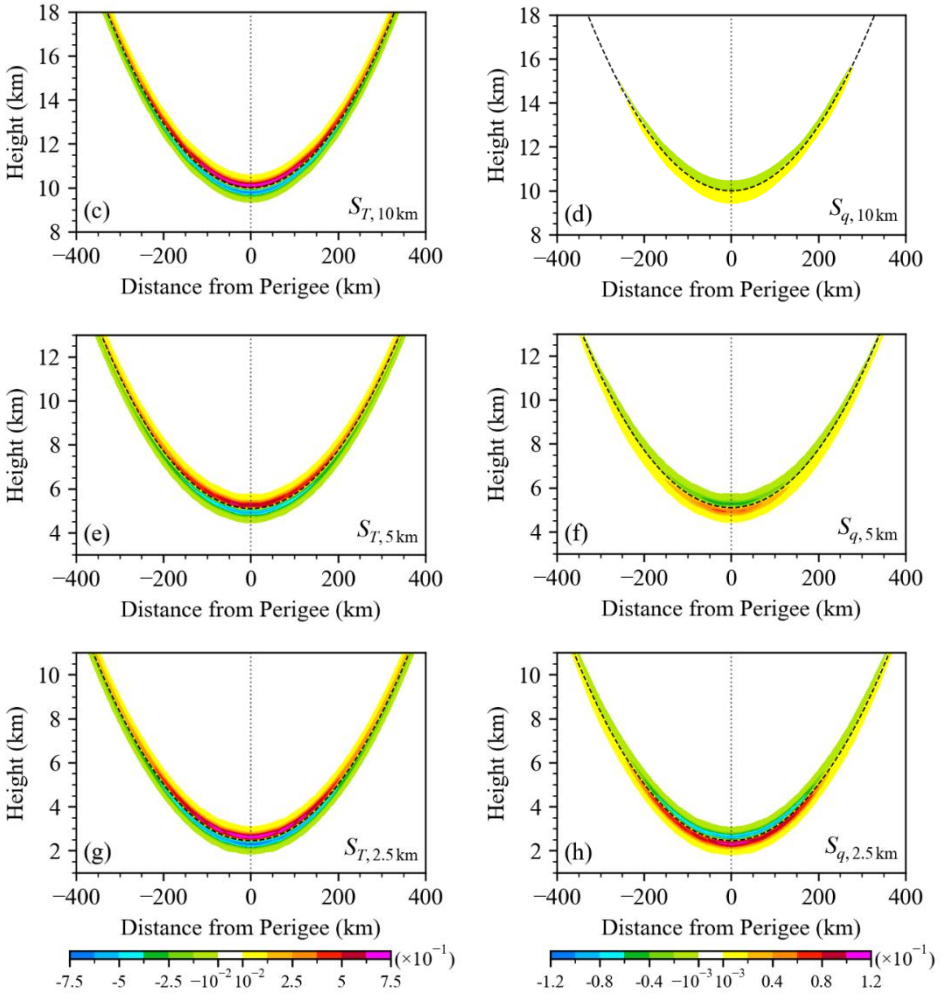


Figure 9. Relative sensitivity of raytracing-derived bending angle (α_{Ray}) to temperature (left panels) and specific humidity (right panels) for RO1 at heights of (a)–(b) 15 km, (c)–(d) 10 km, (e)–(f) 5 km, and (g)–(h) 2.5 km. The rays propagate from the left (GPS side) to right (LEO side).

These height-dependent patterns are closely related to the distributions of refractivity and its gradients along the ray paths. Figure 10 shows that refractivity increases substantially with decreasing target height, while the vertical refractivity gradient at lower levels exhibits both larger magnitudes and more complex local variations. The vertical gradient near 2.5 km exhibits multiple local peaks, whereas gradients near 10 and 15 km are comparatively smooth. Meanwhile, the horizontal refractivity gradients have opposite signs on either side of the tangent point, with larger magnitudes at lower levels. These features indicate that bending-angle sensitivities to temperature and specific humidity depends not only on the local refractivity magnitude but also on the vertical and horizontal gradient structures along the ray path, further highlighting the nonlocal nature of the ray-tracing operator.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

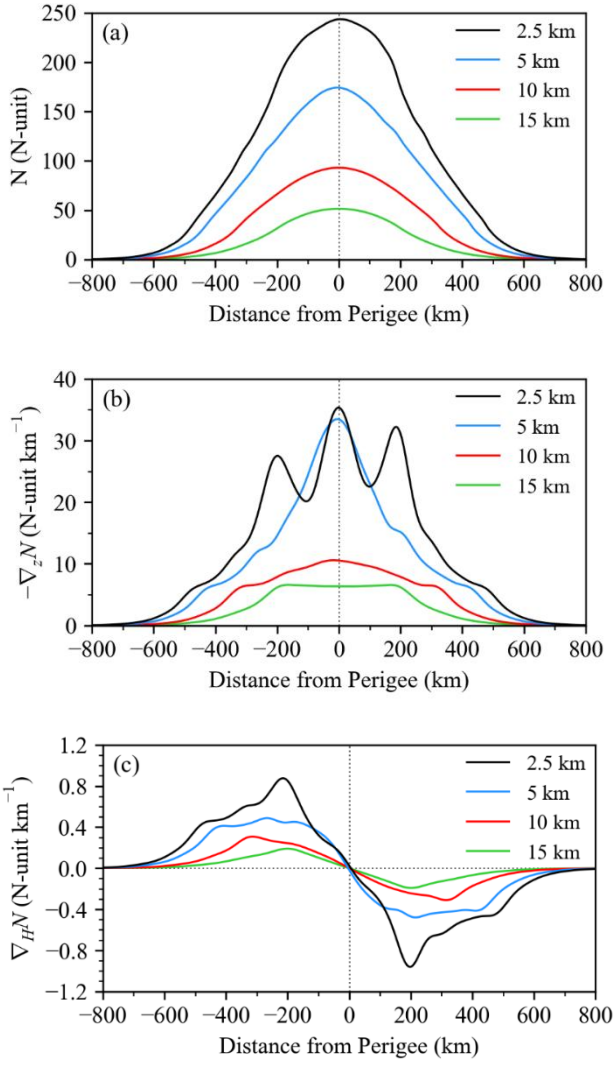


Figure 10. Along-track variations of (a) refractivity, (b) vertical refractivity gradient, and (c) horizontal refractivity gradient for ROs at heights of 2.5, 5, 10 and 15 km.

Figure 11 shows the corresponding relative sensitivities to temperature and specific humidity computed using the AI model. Overall, the AI sensitivities are also distributed along the ray paths, and their spatial distributions and height-dependent variations are broadly consistent with those of the ray-tracing operator. The temperature sensitivity is concentrated near the tangent point and along the curved ray paths on either side, exhibiting vertically adjacent regions of positive and negative sensitivity. The specific-humidity sensitivity is weaker at higher levels but strengthens as the target height decreases, with the spatial extent of the low-level response increasing accordingly. These results indicate that the AI model not only approximates the bending angles simulated by ray tracing but also captures key features of the nonlocal sensitivity of bending angles to temperature and specific-humidity perturbations along the propagation path.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

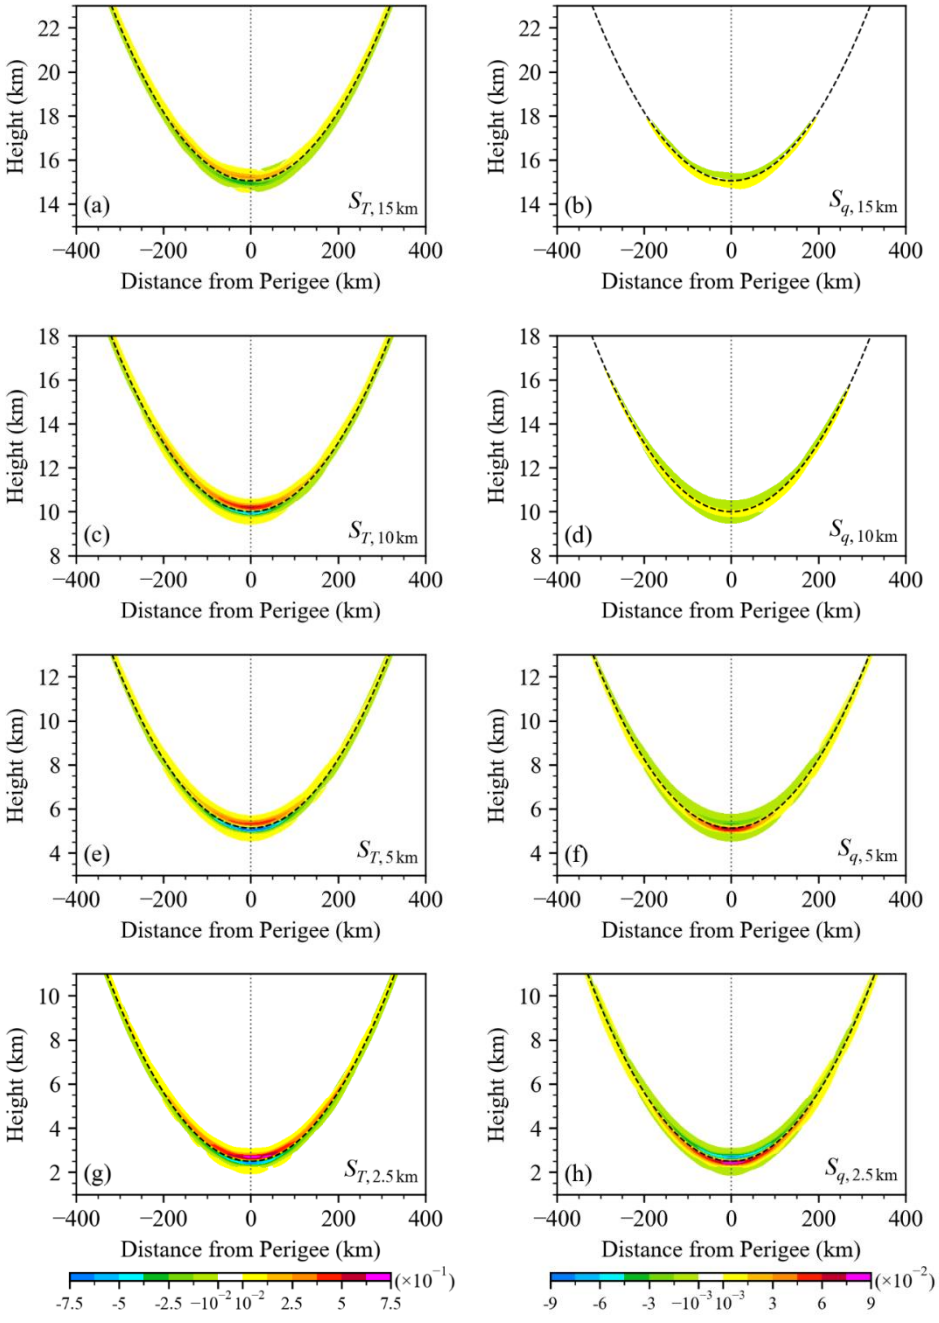


Figure 11. Relative sensitivity of AI-derived bending angle (α_{AI}) to temperature (left panels) and specific humidity (right panels) for RO1 at heights of (a)–(b) 15 km, (c)–(d) 10 km, (e)–(f) 5 km, and (g)–(h) 2.5 km. The rays propogate from the left (GPS side) to right (LEO side).

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

4.5 Computational Efficiency

Computational efficiency is another important criterion for evaluating the AI model as a surrogate for the ray-tracing operator. The tests were conducted on a single CPU core on a computing node equipped with an Intel Xeon E5-2682 v4 processor. The results show that simulating one RO profile takes approximately 9–14 s with ray tracing and 1–3 s with AI inference. For the corresponding adjoint calculations, ray tracing requires approximately 18–22 s per RO profile, whereas the AI model requires approximately 11–15 s. Overall, the AI model is approximately 1.5–3 times faster than ray tracing for the combined forward and adjoint calculations. These results indicate that the AI model can reduce the computational time required for both forward and adjoint calculations, offering a promising approach to improving the computational efficiency of the nonlocal bending-angle observation operator. Further acceleration may be achievable through graphics processing units (GPUs) implementation, particularly when processing multiple RO profiles in parallel.

The reported timings include only the forward and corresponding adjoint calculations, excluding data preprocessing. Therefore, the current results primarily demonstrate the computational advantage of the AI model under the specified test conditions and do not necessarily reflect the speedup achievable with the complete observation operator in practical data assimilation applications.

5 Summary and Conclusions

This study develops an AI-based GNSS RO bending-angle observation operator based on a Transformer encoder and target-dependent multi-head path aggregation for use in data assimilation. The model learns the nonlocal mapping between atmospheric states sampled along a finite ray path and bending angles computed by ray tracing. A composite loss function combining low-level weighted absolute error, piecewise relative error at middle and upper levels, and layer-averaged bias constraints is used to improve simulation accuracy across different heights.

Case studies and independent tests show that the AI model reproduces the main vertical structure of bending angles simulated by ray tracing. Compared with the Abel operator, the AI model achieves higher simulation accuracy and smaller error dispersion at some low-level heights. At middle and upper levels, both operators exhibit high simulation accuracy, while the AI model generally exhibits smaller error dispersion. Incorporating tropopause height information reduces the systematic bias of the AI model near an impact height of 18 km. The relative sensitivity analyses for temperature and specific humidity indicate that the AI model reproduces key features of the nonlocal response patterns of the ray-tracing operator along the propagation path, with generally consistent spatial structures and ranges of sensitivity magnitudes, although some local discrepancies remain.

Computational efficiency tests indicate that the AI model is approximately 1.5–3.0 times faster than ray tracing for the combined forward and adjoint calculations. However, the current timing results exclude data preprocessing and do not represent the end-to-end speedup achievable in an operational data assimilation environment.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Future work will integrate the AI bending-angle observation operator into a numerical weather prediction data assimilation system and conduct experiments assimilating actual GNSS RO observations. Its performance will be compared with that of the conventional ray-tracing and Abel-based observation operators, with particular emphasis on observation fits, analysis increments, assimilation convergence, and forecast impacts. The numerical stability and computational efficiency of the AI observation operator and its adjoint will also be evaluated within the data assimilation system.

Code and data availability

The GNSS RO observations from COSMIC-2 were obtained from the COSMIC Data Analysis and Archive Center (CDAAC, https://cdaac-www.cosmic.ucar.edu/cdaac/). The ERA5 reanalysis data are available from the European Centre for Medium-Range Weather Forecasts (ECMWF) through the Copernicus Climate Data Store (CDS, https://cds.climate.copernicus.eu/). The AI model training code, trained model weights, and data required to reproduce the figures are not publicly archived at present. They may be obtained from the corresponding author upon reasonable request.

Author contributions

XX developed the AI model, conducted the experiments and analyses, and wrote the original draft of the manuscript. XZ contributed to the study design, supervised the research, and reviewed and edited the manuscript. Both authors contributed to the interpretation of the results and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Disclaimer

Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Acknowledgements

We acknowledge the use of computational resources for the AI model training and ray-tracing simulations. We also thank the providers of the ERA5 reanalysis data and COSMIC-2 radio occultation observations for making these datasets available for research.

Financial support

This research was supported by the National Natural Science Foundation of China (Grant No. 42605198).

Review statement

The review statement will be added by Copernicus Publications listing the handling editor as well as all contributing referees according to their status anonymous or identified.

References

Babaud, J., Witkin, A. P., Baudin, M., and Duda, R. O.: Uniqueness of the Gaussian kernel for scale-space filtering, IEEE Trans. Pattern Anal. Mach. Intell., 8, 26–33, https://doi.org/10.1109/TPAMI.1986.4767749, 1986.

ECMWF: ERA5: compute pressure and geopotential on model levels, geopotential height and geometric height, Copernicus Knowledge Base, https://confluence.ecmwf.int/pages/viewpage.action?pageId=226500578 (last access: 12 October 2026), 2021.

Fjeldbo, G., Kliore, A. J., and Eshleman, V. R.: The neutral atmosphere of Venus as studied with the Mariner V radio occultation experiments, Astron. J., 76, 123–140, https://doi.org/10.1086/111096, 1971.

Healy, S. B., Jupp, A. M., and Marquardt, C.: Forecast impact experiment with GPS radio occultation measurements, Geophys. Res. Lett., 32, L03804, https://doi.org/10.1029/2004GL020806, 2005.

Healy, S. B. and Thépaut, J.-N.: Assimilation experiments with CHAMP GPS radio occultation measurements, Q. J. Roy. Meteor. Soc., 132, 605–623, https://doi.org/10.1256/qj.04.182, 2006.

Healy, S. B., Eyre, J. R., Hamrud, M., and Thépaut, J.-N.: Assimilating GPS radio occultation measurements with two-dimensional bending angle observation operators, Q. J. Roy. Meteor. Soc., 133, 1213–1227, https://doi.org/10.1002/qj.63, 2007.

Healy, S. B.: Implementation of the ROPP two-dimensional bending angle observation operator in an NWP system, ROM SAF Report No. 19, EUMETSAT, https://rom-saf.eumetsat.int/general-documents/rsr/rsr_19.pdf, 2014.

Lagerquist, R., Turner, D., Ebert-Uphoff, I., Stewart, J., and Hagerty, V.: Using deep learning to emulate and accelerate a radiative transfer model, J. Atmos. Ocean. Tech., 38, 1673–1696, https://doi.org/10.1175/JTECH-D-21-0007.1, 2021.

https://doi.org/10.5194/egusphere-2026-6771
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Liang, J., Terasaki, K., and Miyoshi, T.: A machine learning approach to the observation operator for satellite radiance data assimilation, J. Meteorol. Soc. Jpn., 101, 79–95, https://doi.org/10.2151/jmsj.2023-005, 2023.

Liang, X., Garrett, K., Liu, Q., Maddy, E. S., Ide, K., and Boukabara, S.: A deep-learning-based microwave radiative transfer emulator for data assimilation and remote sensing, IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens., 15, 8819–8833, https://doi.org/10.1109/JSTARS.2022.3210491, 2022.

Liu, H., Zou, X., Shao, H., Anthes, R. A., Chang, J. C., Tseng, J.-H., and Wang, B.: Impact of 837 GPS/MET bending angle profiles on assimilation and forecasts for the period June 20–30, 1995, J. Geophys. Res.-Atmos., 106, 31771–31786, https://doi.org/10.1029/2001JD000345, 2001.

Shao, H. and Zou, X.: The impact of observational weighting on the assimilation of GPS/MET bending angle, J. Geophys. Res.-Atmos., 107, 4717, https://doi.org/10.1029/2001JD001552, 2002.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I.: Attention is all you need, in: Advances in Neural Information Processing Systems 30, edited by: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S. V. N., and Garnett, R., Curran Associates, Inc., 5998–6008, https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html, 2017.

Zou, X., Vandenberghe, F., Wang, B., Gorbunov, M. E., Kuo, Y.-H., Sokolovskiy, S., Chang, J. C., Sela, J. G., and Anthes, R. A.: A ray-tracing operator and its adjoint for the use of GPS/MET refraction angle measurements, J. Geophys. Res.-Atmos., 104, 22301–22318, https://doi.org/10.1029/1999JD900450, 1999.

Zou, X., Wang, B., Liu, H., Anthes, R. A., Matsumura, T., and Zhu, Y.-J.: Use of GPS/MET refraction angles in three-dimensional variational analysis, Q. J. Roy. Meteor. Soc., 126, 3013–3040, https://doi.org/10.1002/qj.49712657003, 2000.

Zou, X., Liu, H., Anthes, R. A., Shao, H., Chang, J. C., and Zhu, Y.-J.: Impact of CHAMP radio occultation observations on global analysis and forecasts in the absence of AMSU radiance data, J. Meteorol. Soc. Jpn., 82, 533–549, https://doi.org/10.2151/jmsj.2004.533, 2004.

Zou, X.: Adjoint model, in: Atmospheric Satellite Observations: Variation Assimilation and Quality Assurance, Academic Press, 77–96, ISBN 978-0-12-820950-9, 2020.