https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

PCOD-NN v1.0: a differentiable, instrument-independent emulator of monochromatic optical depth for infrared radiative transfer

Junyu Yan[1,2], Chao Liu[1,2,*], Yujia Sun[1,2], Di Di[1,2], and Peng Zhang[3]

[1] State Key Laboratory of Climate System Prediction and Risk Management, School of Remote Sensing and Geomatics Engineering, Nanjing University of Information Science and Technology, Nanjing 210044, China
[2] Collaborative Innovation Center on Forecast and Evaluation of Meteorological Disasters, School of Atmospheric Physics, Nanjing University of Information Science and Technology, Nanjing 210044, China
[3] Meteorological Observation Centre, China Meteorological Administration, Beijing 100081, China

Correspondence to: Chao Liu (chao_liu@nuist.edu.cn)

Abstract. High-spectral-resolution atmospheric modelling and physical retrieval need fast, differentiable infrared radiative transfer to support repeated forward simulations and accurate Jacobians at large scale. Line-by-line models are the accuracy reference but prohibitively slow for practical applications. We introduce this PCOD-NN (principal-component optical-depth neural network), a neural-network-based emulator that reproduces the layer-by-layer, gas-by-gas monochromatic optical depth (OD) over 600–2600 cm^{-1} at 0.01 cm^{-1} resolution for seven gases (H_2O, CO_2, O_3, N_2O, CO, CH_4, and O_2). Instead of fitting radiance of a particular instrument, PCOD-NN predicts OD itself—a quantity upstream of radiative transfer (RT) and instrument convolution—so the same representation couples to different RT solvers, instrument line shapes, and spectral response functions. For each gas, ln(OD) is compressed onto a 500-component PCA basis whose scores are predicted by a residual multilayer perceptron, with separate high- and low-pressure models over 1–1100 hPa. On held-out test sets, the mean relative OD error is less than 0.2% for each of the seven gases. For 90-layer profiles, the V100 GPU implementation is approximately 620–970 times as fast as single-core LBLRTM in the full-resolution seven-gas OD benchmark. The predicted OD spectra are then used in a clear-sky radiative-transfer and instrument-convolution pipeline executed on the GPU to calculate channel brightness temperatures. For the FY-4B/GIIRS hyperspectral infrared sounder, agreement between simulated and observed brightness temperatures is overall comparable to that achieved by RTTOV v14. Without retraining the OD emulator, the framework also simulates infrared channels of the FY-4B/AGRI multispectral imager in a clear-sky ocean scene, achieving agreement with observations broadly comparable to that achieved by RTTOV v14 and demonstrating cross-instrument reuse. Across 12 diverse atmospheric profiles, temperature and water-vapor brightness-temperature Jacobians obtained by automatic differentiation through the forward model have median relative errors of approximately 0.2% and 0.3%, respectively, relative to finite-difference references computed using LBLRTM. By combining accurate emulation of high-resolution, gas-resolved optical depth with differentiable radiative transfer, PCOD-NN provides a common framework for hyperspectral sounding, broadband imaging, and atmospheric sensitivity analysis.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

1 Introduction

Infrared hyperspectral sounders and imagers support many quantitative applications (Menzel et al., 2018), such as temperature and humidity profile retrieval, trace-gas monitoring, data assimilation, and channel radiance simulation. All of these applications rely on a forward operator that maps the atmospheric state to top-of-atmosphere radiance or brightness temperature. The best achievable accuracy of this operator is set by the spectroscopic description of molecular absorption, whose most rigorous implementation is the line-by-line radiative transfer model, such as LBLRTM (Clough et al., 2005). LBLRTM integrates spectral lines one by one and treats continuum absorption explicitly (e.g., the MT_CKD continuum model; Mlawer et al., 2012). In the infrared window, its simulated brightness temperatures agree with high-quality observations to within ~1 K (Alvarado et al., 2013; Turner et al., 2004), so it is often used as the reference truth for other fast models. However, it is computationally expensive: the line-by-line optical-depth computation for one 90-layer profile over the full 2000 cm^{-1} band (at 0.01 cm^{-1} resolution) takes tens of seconds on a single core. LBL simulations are too expensive for high-accuracy, high-resolution quantitative applications, where finer spatial/temporal sampling and higher spectral resolution greatly increase the number of pixels and wavenumbers to be simulated.

Several approaches have been developed to speed up this forward operator. The correlated-k method sorts and integrates absorption coefficients within a channel band, trading the monochromatic spectral structure for speed (Lacis and Oinas, 1991). Fast radiative transfer models such as RTTOV and CRTM (Saunders et al., 2018; Ding et al., 2011) work channel by channel and use regression coefficients to pre-fit the transmittance profile, and are very fast. However, their coefficients are generated in advance for each channel of a specific instrument and are tied to that instrument's channels and spectral response function (SRF), so changing the instrument or the SRF requires regenerating the coefficients. More recently, machine-learning methods have been used to predict atmospheric transmittance and radiance. Gradient boosting trees have been used to predict clear-sky level-to-space transmittance (Liang et al., 2020). A common neural-network design, especially for infrared sounders and imagers, is to directly regress the channel radiance or brightness temperature of one specific instrument (Stegmann et al., 2022; Li et al., 2025).

Among these approaches, one mature and important line is principal-component (PCA/EOF) spectral compression, which reconstructs the spectrum from a small number of low-dimensional scores. PCRTM (Liu et al., 2006) applies PCA to the radiance spectrum. It performs monochromatic radiative transfer at a few selected wavenumbers and then predicts the PC scores from these monochromatic radiances by linear regression. It is used operationally for AIRS, IASI, and CrIS. PC-RTTOV (Matricardi, 2010) belongs to the same family; its predictors are a set of channel radiances computed by the RTTOV fast transmittance model. HT-FRTC (Havemann et al., 2018) performs monochromatic radiative transfer at a few frequencies selected by k-means, but predicts the PC scores with nonlinear kernel regression or a Gaussian process. It is also sensor-independent. These three methods share a common point: the PC scores are all obtained by regression from a set of pre-computed monochromatic or channel radiances, and the prediction loop still contains one reduced radiative transfer computation. In every case the emulation target is the radiance spectrum. Another related line applies PCA directly to optical

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

depth or to inherent optical properties (Natraj et al., 2005, 2010; Spurr et al., 2013, with analytic Jacobians; SDCOMP, Liu et al., 2020; and recent thermal-infrared work by Natraj and Spurr, 2025). These methods usually take the logarithm to compress the dynamic range of OD, which spans several orders of magnitude. This is the same reason behind our ln-OD choice. Neural networks have also been used to predict gas optical properties directly. For example, RRTMGP-NN (Ukkonen et al., 2020; Ukkonen, 2022; Veerman et al., 2021) regresses gas absorption optical depth directly from temperature, pressure, and composition, but without spectral compression.

We choose the layer-by-layer, gas-by-gas monochromatic optical depth as the emulation target to retain the accuracy of line-by-line radiative transfer while remaining flexible across gases, radiative-transfer solvers, and instruments. First, OD is a fundamental physical quantity of the radiative transfer equation and is independent of any instrument. The same OD can be combined with any compatible RT scheme/model, any instrument line shape (ILS), and any SRF within the covered spectral range, so one training serves almost any instruments. Second, keeping the monochromatic spectral structure at 0.01 cm^{-1} resolution lets the emulator stay physically interpretable for even narrower lines, line wings, and weak absorption bands, and it supports almost any band averaging afterwards. Third, the gas-by-gas decomposition lets any mixed atmosphere with different concentration ratios be obtained by combining the single-gas outputs, and it supports clean per-gas attribution.

Building on this formulation, we introduce a principal-component optical-depth neural network (PCOD-NN). PCOD-NN applies PCA compression to gas-by-gas logarithmic optical-depth (ln-OD) spectra. Neural networks predict the corresponding PCA scores directly from atmospheric state variables. The predicted scores are then used to reconstruct high-spectral-resolution monochromatic OD spectra without line-by-line calculations during inference. Because OD is an upstream, instrument-independent physical quantity, the reconstructed spectra can be coupled to different radiative-transfer solvers and instrument response functions. For the cases examined here, PCOD-NN closely reproduces LBLRTM optical depths and instrument-convolved brightness temperatures. It also substantially reduces computational cost and enables accurate Jacobian calculations through automatic differentiation.

The remainder of this paper is organized as follows. Section 2 describes the training data and sampling design. Section 3 presents the PCOD-NN framework, brightness-temperature calculation, and calculation of Jacobians through automatic differentiation. Section 4 evaluates the accuracy of the OD spectra, brightness temperatures, and Jacobians. FY-4B/GIIRS and FY-4B/AGRI are used as examples of a hyperspectral sounder and an imager, respectively. Computational efficiency is benchmarked against LBLRTM for OD calculations and against RTTOV for brightness-temperature calculations. Section 5 discusses the results and their implications.

2 Reference data and sampling design

The training truth is generated by LBLRTM (Clough et al., 2005) in a single-layer configuration. For a given layer-mid temperature T, pressure P, layer half-thickness ratio dP, and target-gas volume mixing ratio (VMR), we compute the monochromatic optical-depth spectrum of that layer over 600–2600 cm^{-1}. This band covers the longwave (LW) and midwave (MW) infrared bands of FY-4B/GIIRS and AGRI (Yang et al., 2017), which is the core range for our target applications. The

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

resolution is 0.01 cm^{-1}. In the single-layer LBLRTM, the layer boundaries are set to $P_{top} = P_{mid}\cdot(1-dP)$ and $P_{bot} = P_{mid}\cdot(1+dP)$, so the actual layer thickness corresponds to $2\cdot P_{mid}\cdot dP$. The training input dP is the half-thickness ratio, $dP = (P_{bot}-P_{top})/(P_{bot}+P_{top}) = \Delta P/(2\cdot P_{mid})$. This convention is used throughout data generation and deployment.

Table 1 overviews the dataset design and parameter ranges, while Figure 1 shows the sampling density and representative per-gas OD spectra of the seven gases. A separate training dataset is generated for each target gas using single-layer LBLRTM calculations. Each training spectrum includes the target-gas absorption together with N_2 background absorption. When the gas-specific predictions are combined, the repeated N_2 background contributions are corrected as described in Sect. 3.2. Pressure spans about three orders of magnitude over 1–1100 hPa. At low pressure, the pressure broadening is small and the absorption lines are very narrow, so a single model cannot handle them together with the high-pressure range. To improve accuracy in the low-pressure range, we split the pressure axis into a high-pressure range and a low-pressure range and model them separately (Table 1). For the high-pressure range (highP, 100–1100 hPa, from the troposphere to the near surface), we draw 50,000 joint samples per gas by Latin hypercube sampling (LHS; McKay et al., 1979), with $T \in [180, 320]$ K and $dP \in [5\times10^{-3}, 0.15]$. The VMR range is given per gas (H_2O 1–4×10^4, CO_2 360–440, O_3 10^{-3}–15, N_2O 0.2–0.5, CO 3×10^{-2}–0.5, CH_4 1–3 ppmv). In the low-pressure range (lowP, 1–100 hPa, the stratosphere), the reasonable range of each parameter changes a lot from one pressure layer to another (for example, O_3 peaks at 1–10 hPa and stratospheric H_2O is very dry). We therefore use a pressure-layer-stratified LHS (60,000 samples per gas), in which each layer is given its own physical range. This avoids wasting samples on physically nonexistent combinations. Here $T \in [175, 295]$ K and $dP \in [5\times10^{-3}, 0.30]$. Sampling is done in the joint space (T, logP, dP, logVMRtarget). P, H_2O, O_3, N_2O, CO, and CH_4 are sampled in log scale, while T, CO_2, and dP are sampled linearly. O_2 is not sampled; its VMR is held fixed at the well-mixed value of 2.09×10^5 ppmv, so only T, P, and dP vary for it. Figure 1a shows the sampling density on the T–P plane; the two ranges join at 100 hPa and each covers its range evenly. Figure 1b shows the per-gas OD spectra of the seven gases for a near-surface layer of a U.S. Standard Atmosphere profile (T = 288 K, P = 1013 hPa, dP = 0.10). The spectra span about 14 orders of magnitude, which illustrates the distribution of the absorption bands and the extreme dynamic range of this problem.

Table 1. Parameter ranges used to generate the high- and low-pressure training datasets.

Parameter	Unit	High-pressure dataset	Low-pressure dataset
T	K	180–320	175–295
P	hPa	100–1.1×10^3	1–100
dP	—	5×10^{-3}–0.15	5×10^{-3}–0.30
H_2O	ppmv	1–4×10^4	1–12
CO_2	ppmv	360–440	350–450
O_3	ppmv	1×10^{-3}–15	5×10^{-2}–20
N_2O	ppmv	0.2–0.5	1×10^{-3}–0.5
CO	ppmv	3×10^{-2}–0.5	4×10^{-2}–0.5
CH_4	ppmv	1–3	8×10^{-2}–2.5
O_2	ppmv	2.09×10^5	2.09×10^5

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

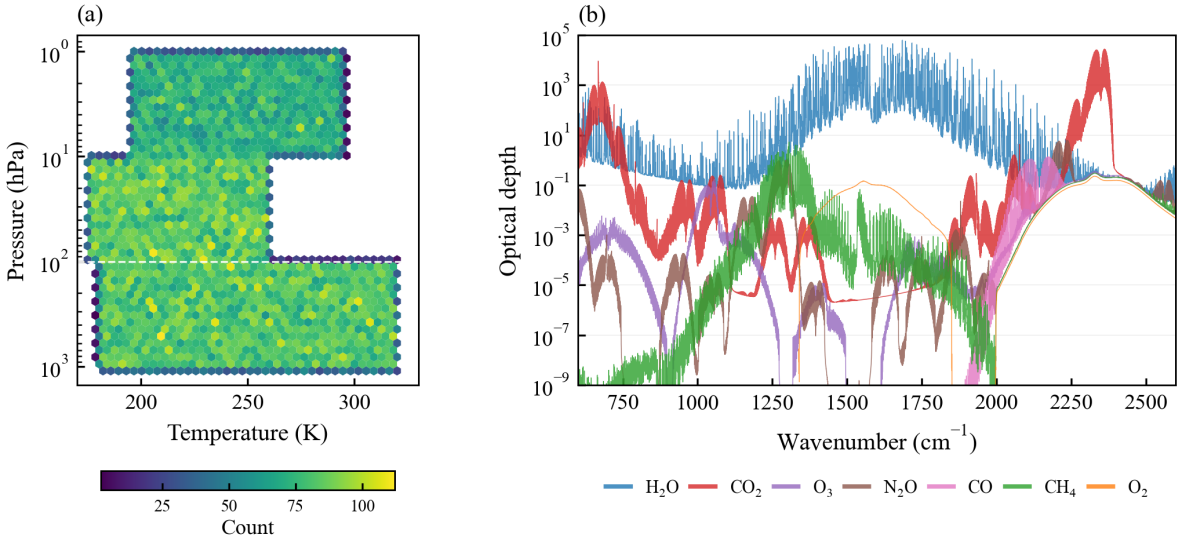


Figure 1. Dataset design. (a) Sampling density on the T–P plane; (b) per-gas OD spectra of the seven gases for a near-surface layer of a U.S. Standard Atmosphere profile (T = 288 K, P = 1013 hPa) ,calculated using wavenumber grid spacings of 0.01 cm⁻¹.

3 Methods

3.1 Optical-depth emulator

The monochromatic OD spectrum is high-dimensional and spans more than ten orders of magnitude. However, for the same gas, the spectral shapes at different (T, P, dP, VMR) are highly correlated and have a clear low-dimensional structure. We apply principal component analysis to ln(OD+ε) (natural logarithm of OD, $\varepsilon = 10^{-13}$) and keep the first k = 500 principal components. For each gas and pressure range, the PCA basis and mean spectrum are fitted using only the training set and are kept fixed during neural-network training and inference. We compress in the ln-OD domain for the following reason. PCA in the linear domain would be dominated by the large OD of strong absorption and would bury the spectral structure of weak absorption. The ln-OD domain instead lets the compression control the relative accuracy of OD evenly across the whole dynamic range.

Figure 2a shows the first five principal-component basis vectors for H_2O in the high-pressure range. The first few PCs capture the broadband envelope, while later PCs gradually add the details of narrow lines and band wings. Figure 2b shows the score matrix of 5000 samples × 500 PCs, with the samples sorted by pressure. The scores are concentrated in the first few components and drop quickly as the index increases. Figure 2c is the key evidence for the best achievable accuracy. With k = 500 reconstruction (10,000 samples × 2000 wavenumber bins, high- and low-pressure combined), the relative error is below 0.1% almost everywhere and better than 0.01% over most of the range. This shows that the PCA truncation itself is far from the main source of error. As long as the network predicts these 500 scores accurately, the overall accuracy is guaranteed. This step turns a 200,000-dimensional spectral regression problem into a 500-dimensional one with a known error bound.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

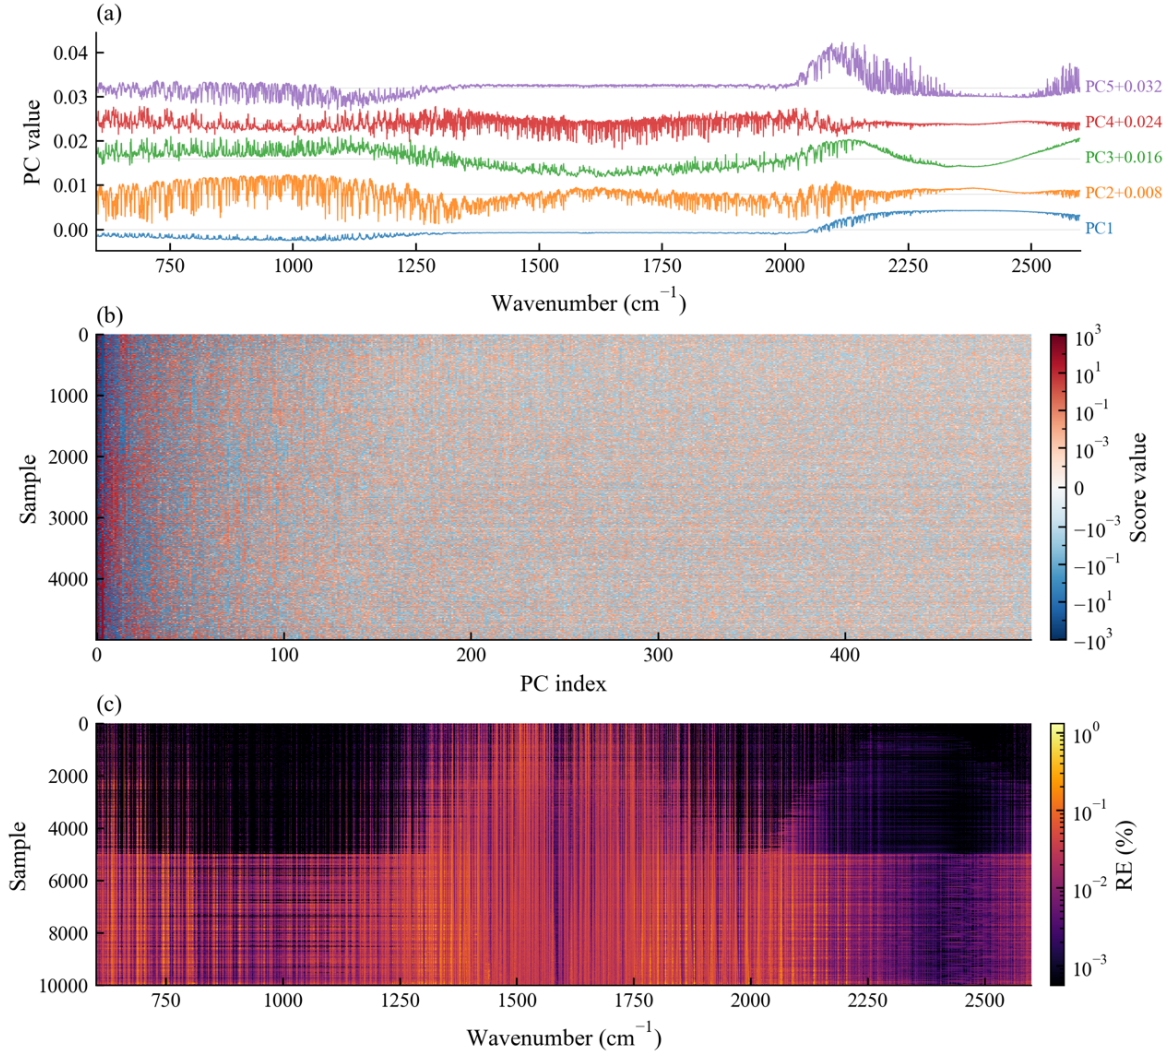


Figure 2. PCA spectral compression. (a) First five principal-component bases for H_2O in the high-pressure range (shifted vertically for clarity); (b) PC score matrix (5000 samples × 500 PCs, samples sorted by pressure); (c) relative error of the k = 500 reconstruction (10,000 samples × 2000 wavenumber bins, high- and low-pressure combined), ≲0.1% everywhere and <0.01% over most of the range, showing that the best achievable accuracy of the PCA method is sufficient.

The emulator is named PCOD-NN (principal-component optical-depth neural network); its principal-component scores are predicted by a residual multilayer perceptron (ResMLP; He et al., 2016). Its input is a small feature vector derived from the layer state (T, P, dP, VMR_{gas}): temperature, log-pressure, layer half-thickness, and log-VMR. For the high-dynamic-range gases (O_3, N_2O, CO, CH_4), two engineered terms—a log column-amount log(P·dP·VMR) and an inverse temperature 1000/T—are appended, giving a six-feature input (four for H_2O, CO_2, and O_2). An input projection maps this feature vector to a 768-dimensional hidden layer, followed by six residual blocks (equivalent to twelve 768-dimensional hidden layers; see Fig. 3). Each block contains LayerNorm, a GELU activation (Hendrycks and Gimpel, 2016), and dropout (0.1). The output is the 500

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

principal-component scores (implemented as 501 outputs, including one global scale term). The predicted scores are multiplied with the PCA bases and summed, and an exponential then recovers the full OD spectrum:

$$\mathrm{OD}(\nu) = \exp\left(\sum_{i=1}^{500} s_i \ \mathrm{PC_i}(\nu) + \mu(\nu) + a\right) - \varepsilon \quad (1)$$

where ν is the wavenumber; s_i (i = 1, …, 500) are the principal-component scores predicted by the ResMLP; $\mathrm{PC_i}(\nu)$ are the corresponding fixed PCA basis vectors; μ(ν) is the mean ln-OD spectrum; a is the predicted global scale term; and $\varepsilon = 10^{-13}$ is the small offset that regularizes the logarithm.

Training uses a spectral loss computed in ln-OD space rather than in score space: the predicted scores are mapped back through the fixed PCA basis to ln OD(ν), and the loss is the mean squared error between the predicted and true ln OD over all wavenumbers. Because the gradient passes through the PCA basis, every score is weighted by its true spectral contribution, which prevents the errors of the later scores from being amplified during reconstruction. Each model has about 7.5×10^6 parameters (about 30 MB in float32), and its PCA basis is about 374 MB.

All fourteen models (seven gases × two pressure ranges) use the same optimization setup. The network inputs and the PCA-score targets are each standardized to zero mean and unit variance. For the three gases with deep, saturated near-zero-OD regions (O_3, N_2O, CO), a floor mask down-weights (weight 0.1) the wavenumbers whose true OD is below 10^{-12}, so the log floor does not dominate the gradient. We optimize with AdamW (weight decay 1×10^{-3}; Loshchilov and Hutter, 2019) and a one-cycle learning-rate schedule (Smith and Topin, 2019; maximum learning rate 1×10^{-4}, 20% warm-up, cosine annealing), a batch size of 128, and gradient clipping at a max-norm of 1.0. Each model is trained for up to 1500 epochs with early stopping (patience 500) on the validation loss. For each gas and pressure range, the samples are split 70/15/15 into training, validation, and test sets with a fixed random seed, and neural-network prediction accuracy is evaluated on the held-out test set. Each model is trained on a single V100 GPU.

For each gas and each pressure range (highP/lowP), we train a separate (PCA basis, ResMLP) pair. At inference, each layer is routed to the model of its pressure range according to the input pressure P. When the P array crosses the range boundary, routing is done sample by sample. This work uses 0.01 $\mathrm{cm^{-1}}$ resolution and covers 1–1100 hPa with the two models (lowP and highP).

3.2 Mixed-atmosphere OD and brightness-temperature simulation

The mixed-atmosphere OD is obtained by summing the seven gas-specific predictions and applying an approximate correction for repeated N_2 background absorption using a lookup table generated from LBLRTM calculations without the target gases. The full brightness-temperature simulation pipeline is: profile levels (P_{lev}, T_{lev}, VMR_{lev}) → layer-mid parameters (P_{mid}, dP, T_{repr}) → per-gas network prediction → mixed-atmosphere OD construction → linear-in-τ analytic layer-source radiative transfer → top-of-atmosphere radiance I_{TOA} → instrument convolution (GIIRS ILS or AGRI SRF) → analytic inverse Planck → channel brightness temperature. The radiative transfer uses the analytic solution in which the in-layer source function varies linearly with optical depth, and the profile arrays are ordered from TOA to the surface. On the instrument side, one can switch

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

between the GIIRS ILS convolution with resampling and the SRF integration in radiance space for each FY-4B/AGRI channel (I_{ch} = Σ I·SRF / Σ SRF), followed by an inverse Planck at the SRF-weighted equivalent central wavenumber.

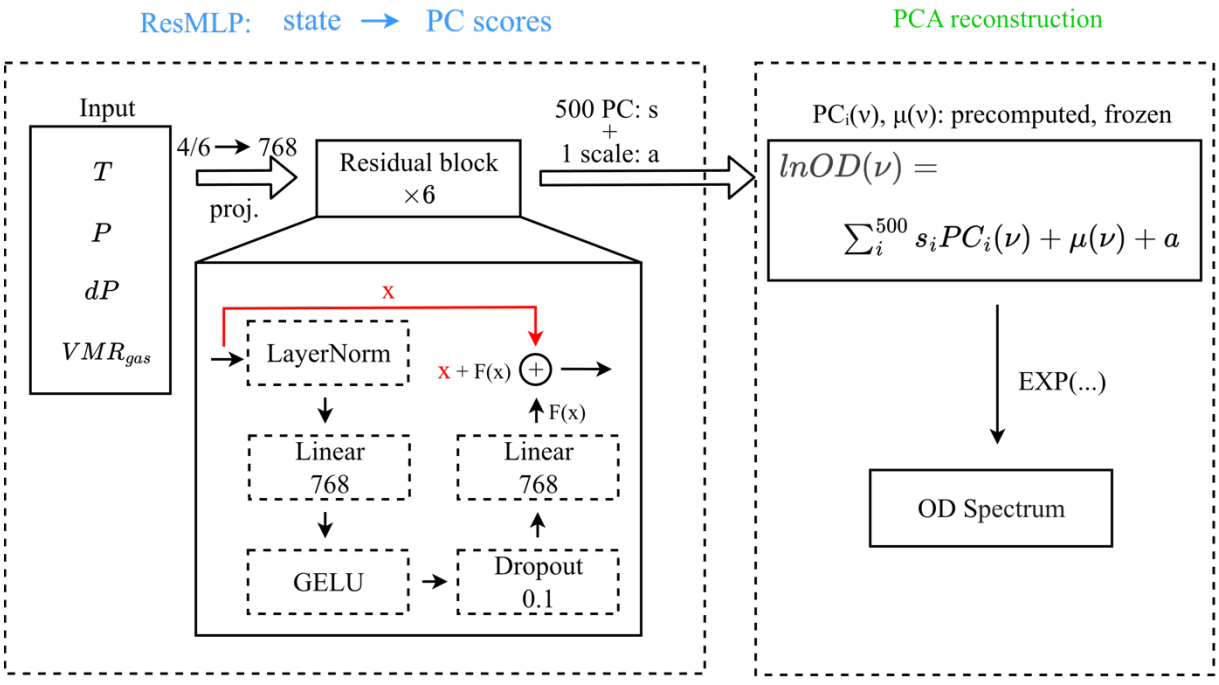


Figure 3. Algorithm framework. Left (ResMLP, state → principal-component scores): the feature vector derived from the layer state (T, P, dP, VMR_{gas})—four features for $H_2O/CO_2/O_2$, six for $O_3/N_2O/CO/CH_4$ (Sect. 3.1)—is linearly projected to 768 dimensions, then passed through six residual blocks (each block: LayerNorm–Linear–GELU–Dropout–Linear, with a residual skip connection x + F(x)), giving 500 principal-component scores s and one global scale term a. Right (PCA reconstruction): with a fixed PCA basis (PC_i and the mean μ are fitted offline and not trained), ln OD(ν) = Σ_i s_i·PC_i(ν) + μ(ν) + a is obtained, and exp (minus ε) recovers the per-gas OD spectrum.

The GPU implementation keeps the whole pipeline (network forward pass, PCA reconstruction, and exponential) in GPU memory. The seven gases are accumulated on the device, and only one transfer back to the host is needed per pressure range. Inference is automatically split along the sample dimension with a fixed block size (512), so the intermediate tensor between GPU and host stays at (512, n_wn). The peak memory is therefore decoupled from the number of input layers: in our tests, the GPU peak is about 4.8 GiB for both 1024 and 16,000 layers (seven gases). Large batches therefore do not overflow memory, and the throughput does not change with the number of layers. The CPU path uses numpy (multithreaded BLAS) and is numerically consistent with the GPU path (relative difference below 0.1%, only from the float32 accumulation order).

3.3 Differentiable forward model and automatic-differentiation Jacobians

For the Jacobian calculations considered here, layer pressures and layer-thickness parameters are held fixed, so each layer remains assigned to the same pressure-specific submodel. Under these conditions, the forward pipeline is differentiable with

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

respect to temperature, water vapor, and the trace-gas concentrations considered in this study: per-gas ResMLP → reconstruction with the fixed PCA basis → exponential → mixed-atmosphere OD construction → radiative transfer → instrument convolution → inverse Planck conversion. The neural networks approximate the mapping from atmospheric state to PCA scores and a global scale term; the subsequent reconstruction and radiative-transfer operations are differentiated as implemented. The Jacobian (K matrix) with respect to the atmospheric state can therefore be obtained directly by automatic differentiation in PyTorch (Baydin et al., 2018; Paszke et al., 2019), without hand-written tangent-linear or adjoint code and without finite-difference perturbations for Jacobian computation (Hatfield et al., 2021). This follows the same idea as PCA-RT acceleration, which requires analytic linearization of the whole PCA enhancement (Spurr et al., 2013); here we leave this to automatic differentiation instead.

Automatic differentiation propagates derivatives through the predicted PCA scores and global scale term, mixed-atmosphere OD construction, radiative transfer, instrument convolution, and inverse Planck conversion. For temperature derivatives, this includes both the temperature dependence of absorption and that of the Planck source function. The resulting channel brightness-temperature Jacobians are evaluated against finite-difference references based on LBLRTM in Sect. 4.6. Derivatives can be propagated through the fixed PCA reconstruction without explicitly constructing the full Jacobian of the monochromatic spectrum with respect to the atmospheric state. The smooth GELU activation supports smooth network outputs and derivatives within each pressure-specific submodel. The monochromatic OD and radiance derivatives are independent of the instrument response. Instrument-specific brightness-temperature Jacobians are obtained by applying the corresponding instrument-convolution and inverse-Planck derivatives.

4 Results

4.1 Single-gas optical-depth accuracy

Figure 4 summarizes the scatter of predicted OD against LBLRTM truth for the seven gases on the 0.01 cm^{-1} test set. The data span about ten orders of magnitude and stay close to the 1:1 line. The relative-error panel below shows that almost all errors fall within ±1%, becoming smaller at large OD and slightly wider at weak absorption (OD $\lesssim 10^{-2}$). The latter is exactly the natural behavior of ln-OD PCA near zero absorption.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

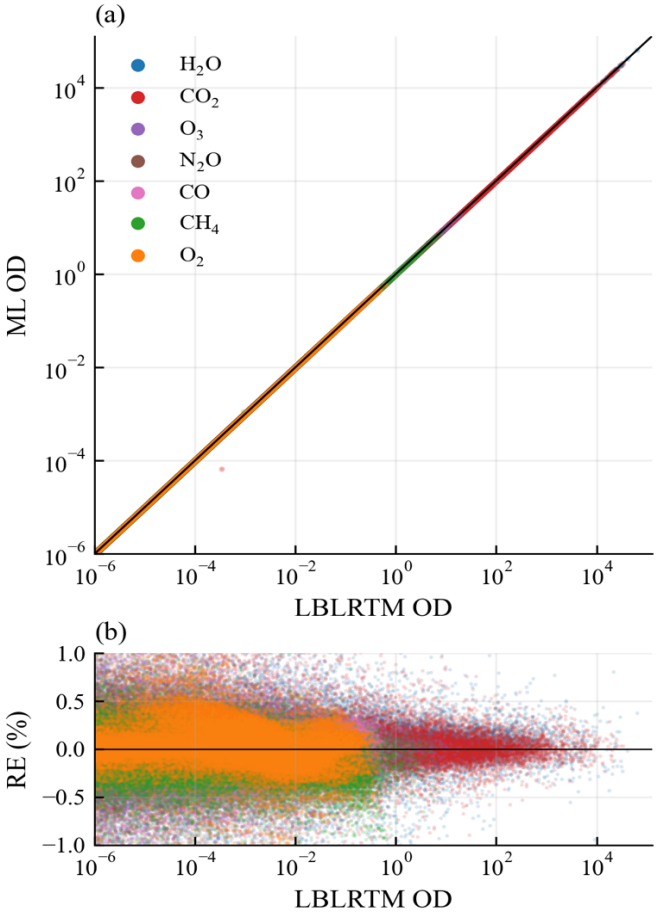


Figure 4. ML-predicted OD versus LBLRTM truth for the seven gases on the 0.01 cm⁻¹ independent neural-network test set. (a) Scatter of ML-predicted OD against LBLRTM truth (black line is the 1:1 reference); (b) distribution of the relative error as a function of LBLRTM truth OD.

Figures 5 and 6 show, gas by gas, the predicted spectra overlaid on the LBLRTM truth (left column) and the relative-error spectra (right column). They correspond, respectively, to a low-pressure case (mid-stratosphere; T = 226 K, P = 30.0 hPa, dP = 0.035) and a high-pressure case (near surface; T = 291 K, P = 950.3 hPa, dP = 0.100). The relative-error spectra are drawn with two normalizations. The gray line is the per-wavenumber relative error (the denominator is the LBLRTM OD at each wavenumber, so the deep spectral valleys are amplified as the denominator approaches zero). The red line uses the in-band peak OD as the denominator, which better reflects the actual effect on radiation and is not amplified in these deep valleys. Figure 7 uses the same convention. In both cases, the predicted spectra of H_2O, CO_2, O_3, N_2O, CO, and CH_4 are almost indistinguishable from the truth over the whole spectrum and the whole dynamic range. The mean absolute per-wavenumber relative error (mean|RE|) is in the range 0.05%–0.16%. The errors mainly appear at strong-gradient locations such as line wings and band edges, as expected.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

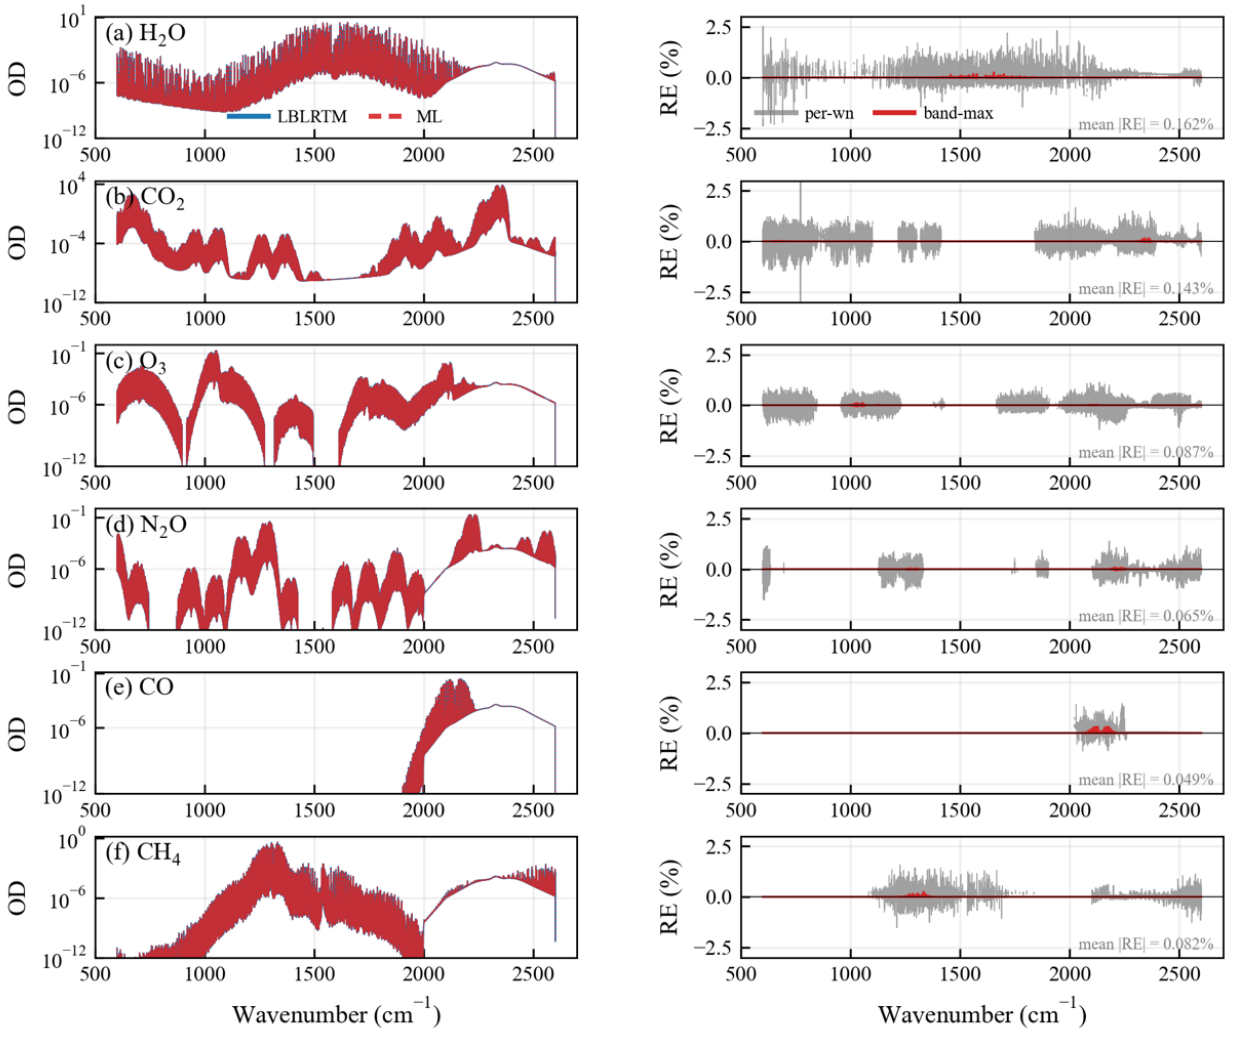


Figure 5. Single-gas low-pressure case (mid-stratosphere; T = 226 K, P = 30.0 hPa, dP = 0.035). (a)–(f) are H_2O, CO_2, O_3, N_2O, CO, and CH_4 in order; for each gas, the left column overlays the predicted spectrum on the LBLRTM truth, and the right column is the relative-error spectrum, with mean|RE| labeled in each error panel.

Table 2 gives the per-gas statistics on the test set. In the high-pressure range (100–1100 hPa), the mean relative error of all seven gases is below 0.2% and the 99th percentile is below 1%. In the low-pressure range (1–100 hPa), the mean relative error is also below 0.2%. The 99th percentiles of CO_2 and O_3 are 2.0% and 1.7%, respectively, with the errors concentrated in the narrow-line region of the upper stratosphere at 1–10 hPa (all gases are clean in the 10–100 hPa segment).

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

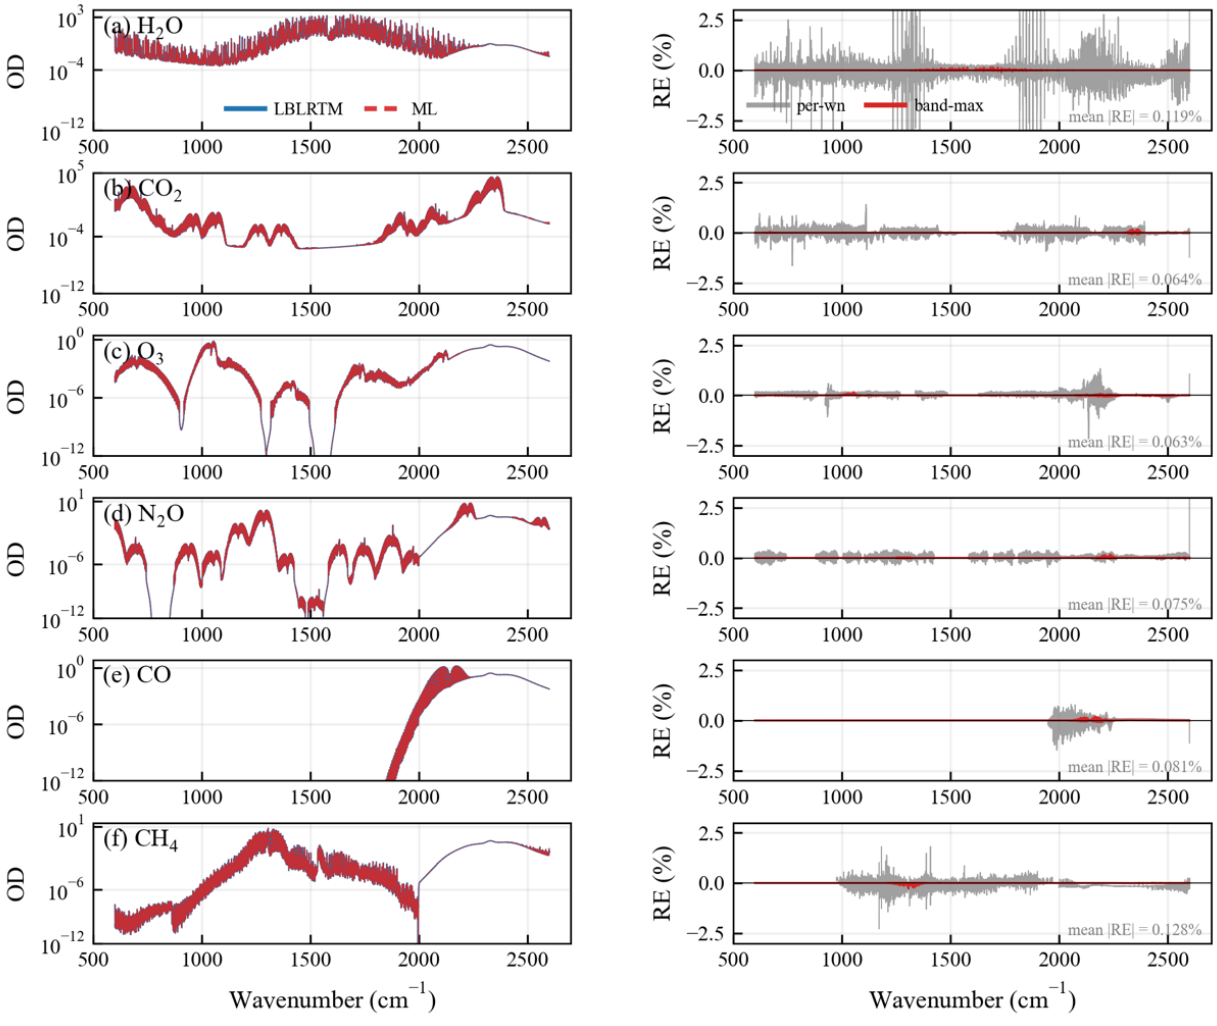


Figure 6. Single-gas high-pressure case (near surface; T = 291 K, P = 950.3 hPa, dP = 0.100); layout as in Figure 5.

4.2 Total optical depth of the seven-gas mixture

Figure 7 compares the total OD (obtained by combining the seven single-gas predictions) with the LBLRTM truth (all seven gases), for two representative cases (a: high-pressure, near surface; b: low-pressure, mid-stratosphere). The top row overlays the total OD spectra, the middle row shows the absolute difference PCOD-NN − LBLRTM, and the bottom row shows the relative-error spectrum. The per-sample mean relative error of the two cases is 0.078% and 0.118%, respectively. The absolute difference shows visible spikes only in regions of very large OD, such as the strong CO_2 band (about 2350 cm^{-1}), but the relative error remains small. The bottom row uses the two normalizations (same as Fig. 5: gray for per-wavenumber, red for in-band maximum) and shows that the error is overall at the ±1% level within the bands.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 2. Single-gas OD relative error on the test set (OD > 10^{-6}).

Gas	highP Mean	highP P99	lowP Mean	lowP P99
H_2O	0.146%	0.927%	0.094%	0.641%
CO_2	0.096%	0.471%	0.199%	2.011%
O_3	0.090%	0.488%	0.181%	1.673%
N_2O	0.064%	0.308%	0.113%	0.802%
CO	0.128%	0.529%	0.059%	0.359%
CH_4	0.158%	0.700%	0.086%	0.548%
O_2	0.162%	0.773%	0.151%	0.636%

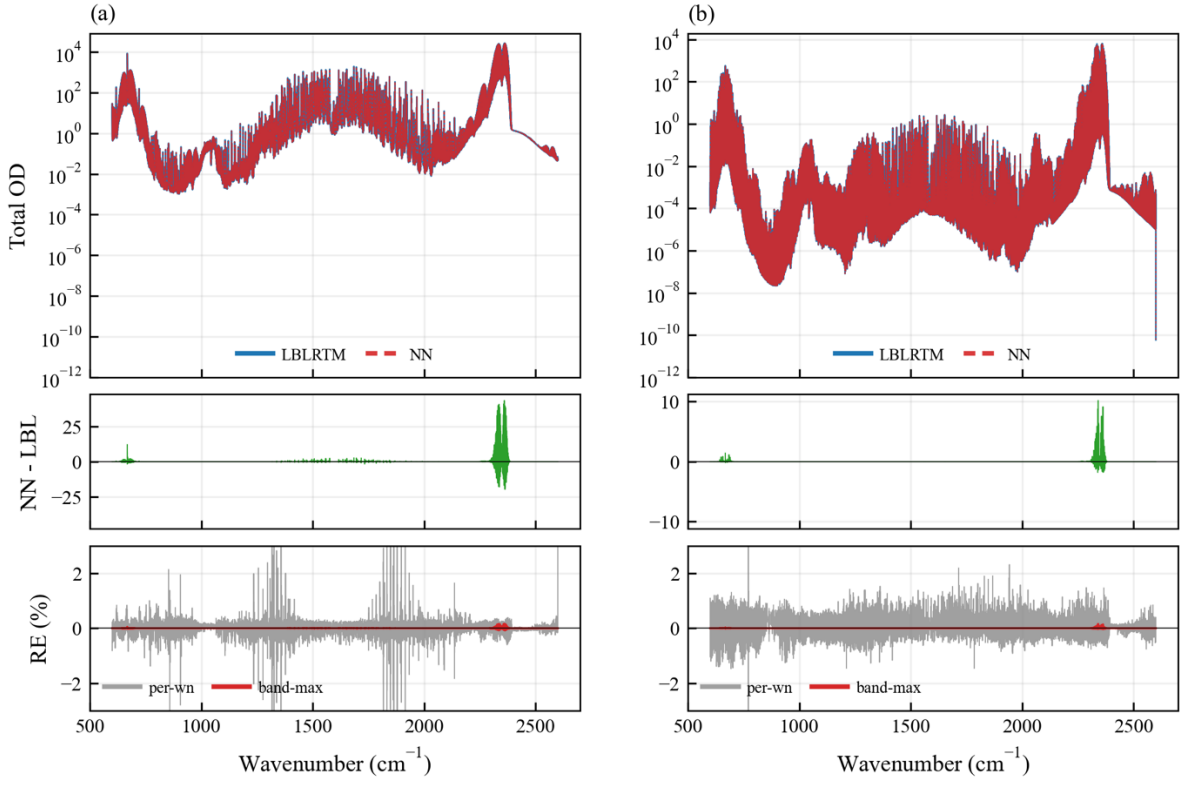


Figure 7. Total OD comparison for the seven-gas mixture (PCOD-NN vs LBLRTM truth). (a) High-pressure, near surface (T = 291 K, P = 950.3 hPa); (b) low-pressure, mid-stratosphere (T = 226 K, P = 30.0 hPa). For each case, the top row overlays the total OD spectra, the middle row shows the PCOD-NN − LBLRTM absolute difference, and the bottom row shows the relative-error spectrum; the per-sample mean relative error is (a) 0.078% and (b) 0.118%.

4.3 Computational performance

Figure 8 compares the wall-clock time required to generate seven-gas monochromatic optical-depth (OD) spectra for a 90-layer profile over 600–2600 cm^{-1} at 0.01 cm^{-1} resolution. The comparison uses OD availability in host memory as its endpoint:

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

GPU timing includes transfer of the computed OD to host memory, whereas LBLRTM runs on the CPU and its output is already available there. Single-core LBLRTM reports approximately 42 s for the line-by-line OD calculation. At a batch size of 16 profiles, PCOD-NN takes approximately 2.31 s per profile on one CPU core, about 18 times faster than LBLRTM. Across the measured batch sizes, PCOD-NN takes approximately 43–68 ms per profile on a V100 GPU, corresponding to a speedup of approximately 620–970 times relative to single-core LBLRTM. At a batch size of 16 profiles, model computation and OD transfer to host memory account for approximately 44% and 56% of the measured GPU runtime, respectively.

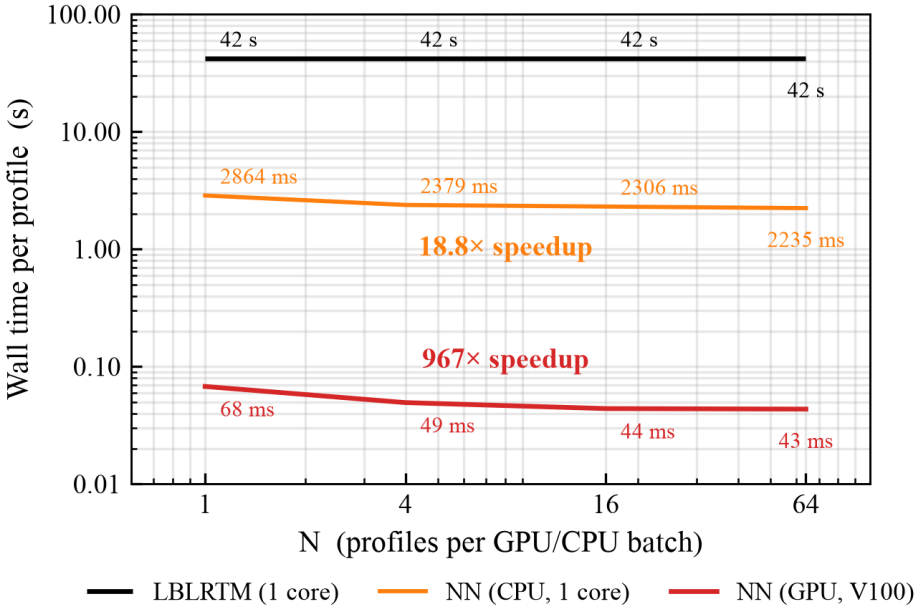


Figure 8. Computational time for seven-gas monochromatic optical-depth (OD) spectra. Per-profile wall-clock time at different batch sizes N is shown for PCOD-NN on one CPU core and a V100 GPU, with single-core LBLRTM as the reference.

When the subsequent radiative-transfer calculation is also performed on the GPU, the computed OD can be used there directly without transferring the full-resolution spectra to host memory. We use this workflow in a separate clear-sky FY-4B/GIIRS benchmark with the same 64 atmospheric profiles for both methods, each containing 90 layers (Table 3). In this benchmark, the measured output is the brightness temperature for 1,682 GIIRS channels. RTTOV v14 computes the instrument channels directly, whereas PCOD-NN evaluates OD at 158,450 monochromatic wavenumbers, performs high-resolution radiative transfer, and convolves the radiances with the GIIRS instrument line shape. The monochromatic OD and radiance arrays remain on the GPU, while the final channel brightness temperatures are transferred to host memory; thus, this benchmark differs from Figure 8 in both spectral grid and output. Timing begins after profile preparation and model initialization. For the GPU implementation, it includes the computation and the transfers needed to deliver the final channel brightness temperatures to host memory. For the common 1,682-channel output, RTTOV required 27.86 ms per profile on one CPU core and 2.69 ms

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

on 16 CPU cores; PCOD-NN required 2,446.85 ms and 1,299.47 ms on the corresponding CPU configurations, and 24.26 ms on a V100 GPU. The V100 runtime was close to that of single-core RTTOV, while 16-core RTTOV was approximately nine times faster. Beyond this timing comparison, the instrument-independent OD representation also supports the GIIRS and AGRI applications and the automatic-differentiation Jacobian analyses presented in Sects. 4.4–4.6.

Table 3. Wall-clock time for clear-sky FY-4B/GIIRS simulations of 64 atmospheric profiles, each with 90 layers. All methods produce brightness temperatures for the same 1,682 channels. Timing begins after profile preparation and model initialization; the GPU timing includes transfer of the final channel output to host memory.

Method	Platform	Total time (s)	Time/profile (ms)	Throughput (profiles s^{-1})
RTTOV v14	CPU, 1 core	1.783	27.86	35.89
RTTOV v14	CPU, 16 cores	0.172	2.69	372.32
PCOD-NN	CPU, 1 core	156.598	2446.85	0.409
PCOD-NN	CPU, 16 cores	83.166	1299.47	0.770
PCOD-NN	NVIDIA V100 GPU	1.553	24.26	41.22

4.4 Validation against FY-4B/GIIRS observations

Figure 9 presents an illustrative single-field-of-view comparison based on an ERA5 atmospheric profile adjusted using collocated GIIRS observations to reduce systematic simulation–observation differences (Hersbach et al., 2020; 24.8°N, 124.0°E, over ocean; 2024-08-02 01:45 UTC; Tskin=303.3 K; VZA = 36.2°). Brightness temperatures calculated using PCOD-NN, RTTOV v14, and LBLRTM are compared with collocated FY-4B/GIIRS observations in the longwave [679, 1131] cm^{-1} and midwave [1649, 2251] cm^{-1} bands. For the statistics, 6 cm^{-1} is excluded from both ends of each band, leaving 705 longwave and 945 midwave channels. Observational residuals are defined as simulated minus observed brightness temperature; μ and σ denote the mean and standard deviation. The PCOD-NN−LBLRTM curve shows the difference between the brightness temperatures calculated using the two methods. In the longwave band, PCOD-NN − OBS gives μ = +0.27 K, σ = 0.46 K, and RMSE = 0.53 K, while RTTOV − OBS gives μ = −0.07 K, σ = 0.47 K, and RMSE = 0.47 K. In the midwave band, PCOD-NN − OBS gives μ = +0.07 K, σ = 0.78 K, and RMSE = 0.78 K, compared with μ = −0.03 K, σ = 0.77 K, and RMSE = 0.77 K for RTTOV − OBS. The direct PCOD-NN − RTTOV differences are μ = +0.34 K and σ = 0.50 K in the longwave band, and μ = +0.10 K and σ = 0.35 K in the midwave band. For this field of view, PCOD-NN and RTTOV v14 show comparable agreement with GIIRS observations. The PCOD-NN−LBLRTM brightness-temperature differences are very small across most channels, with larger departures near the high-wavenumber end of the midwave band. This close agreement shows that PCOD-

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

NN, trained on LBLRTM-generated optical-depth spectra, closely reproduces the channel brightness temperatures obtained using LBLRTM for this field of view.

We then evaluated PCOD-NN and RTTOV v14 using a separate, multi-scene sample of clear-sky ocean fields of view over the western North Pacific. For this evaluation, the ERA5 atmospheric profiles were used without adjustment to the GIIRS observations. The sample covers FY-4B/GIIRS acquisitions in April, July, August, October, and December 2024. Following strict AGRI cloud screening, fields of view were retained only when the Effect Score in both the GIIRS longwave and midwave quality-assessment matrices was 100. This yielded 426 fields of view from nine acquisition times (11 scan files), comprising 220 daytime and 206 nighttime samples (Fig. 10). For each selected field of view, ERA5 pressure-level temperature, specific humidity, and ozone were matched to its GIIRS geolocation by bilinear spatial interpolation and linear interpolation between the hourly ERA5 fields bracketing the scan midpoint. The matched fields were mapped to a common radiative-transfer pressure grid, and the resulting atmospheric profile was supplied to both PCOD-NN and RTTOV v14. Surface pressure and surface temperature were obtained from the corresponding ERA5 single-level fields. The same 6 cm^{-1} band-edge exclusion retained 705 longwave and 945 midwave channels per field of view; the statistics below are pooled over the matched field-of-view–channel pairs.

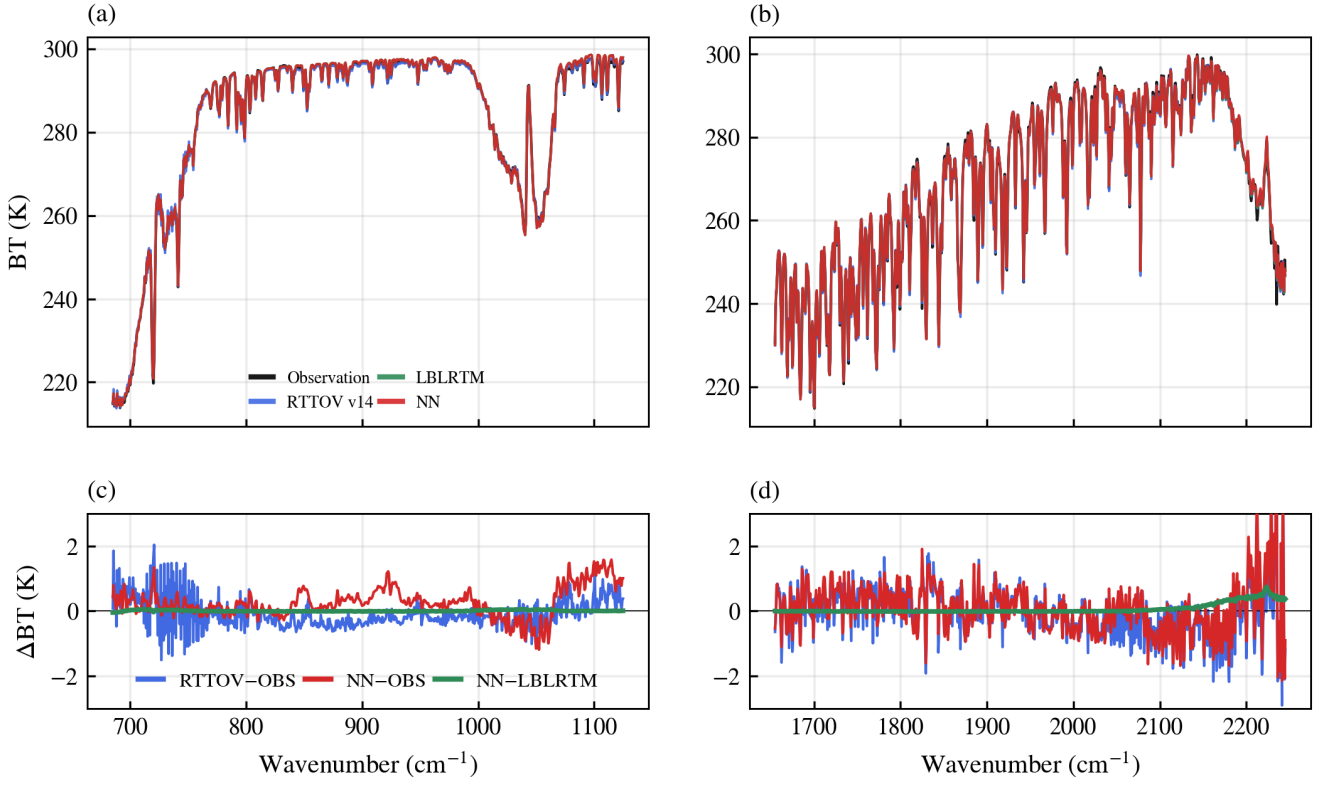


Figure 9. Brightness-temperature validation (24.8°N, 124.0°E, over ocean; T_{skin} = 303.3 K, VZA = 36.2°). (a, b) GIIRS observations and simulations from PCOD-NN (NN), RTTOV v14, and LBLRTM in the longwave and midwave bands, respectively. (c, d) Brightness-temperature differences: RTTOV−OBS, NN−OBS, and NN−LBLRTM. OBS denotes GIIRS observations.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

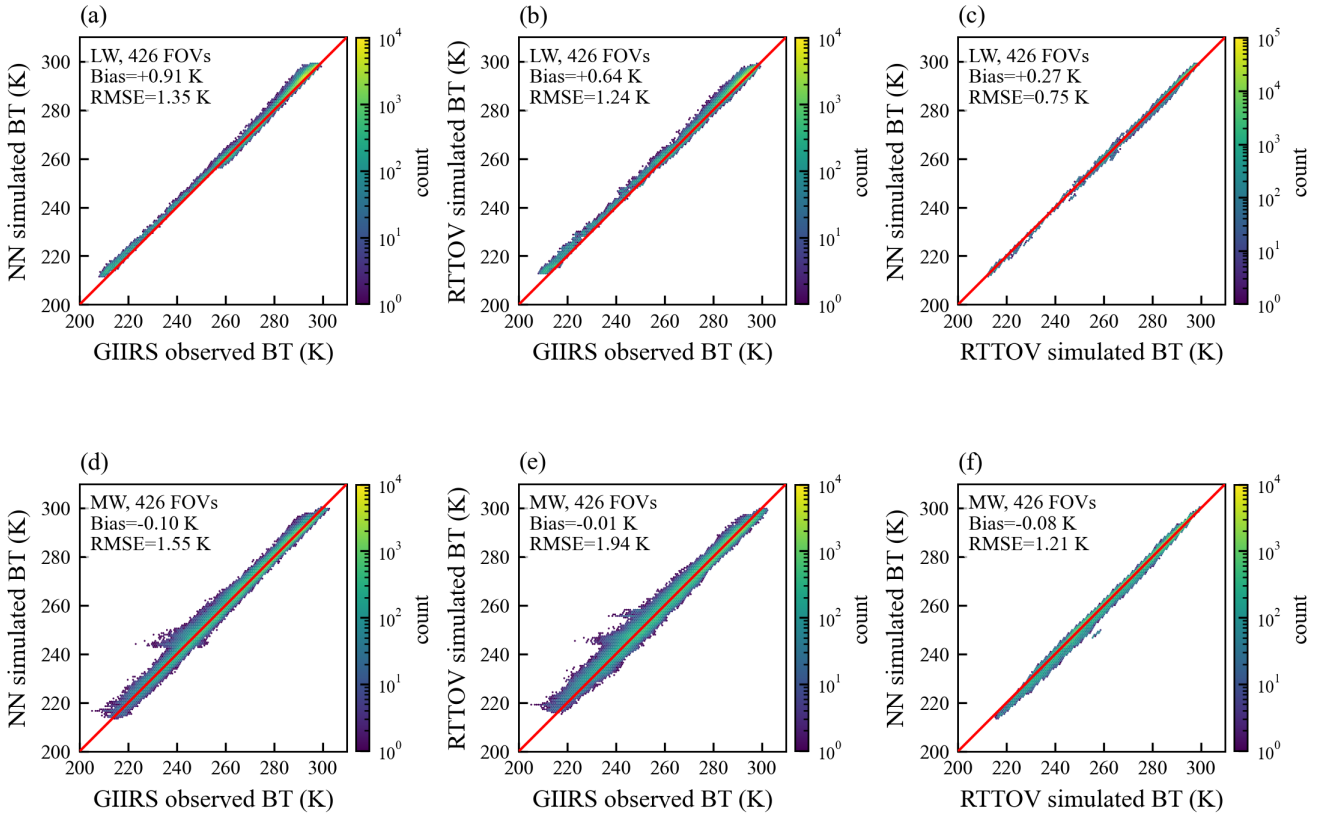


Figure 10. Clear-sky ocean brightness-temperature comparisons for 426 FY-4B/GIIRS fields of view from nine acquisition times (11 scan files) in April, July, August, October, and December 2024. The samples include 220 daytime and 206 nighttime fields of view over the western North Pacific (approximately 17–21°N, 122–129°E). All fields of view pass strict AGRI clear-sky screening and have a GIIRS Effect Score of 100 in both bands. The top and bottom rows show the longwave and midwave bands, respectively. The columns compare PCOD-NN with GIIRS (a, d), RTTOV v14 with GIIRS (b, e), and PCOD-NN with RTTOV v14 (c, f). Colour indicates the local number of field-of-view–channel pairs on a logarithmic scale; red lines indicate one-to-one agreement. Bias is defined as the ordinate minus the abscissa.

Relative to GIIRS, PCOD-NN has a bias and RMSE of +0.91 and 1.35 K in the longwave band and −0.10 and 1.55 K in the midwave band (Fig. 10a, d). The corresponding RTTOV v14 values are +0.64 and 1.24 K, and −0.01 and 1.94 K (Fig. 10b, e). The direct PCOD-NN minus RTTOV v14 comparison gives a bias and RMSE of +0.27 and 0.75 K in the longwave band and −0.08 and 1.21 K in the midwave band (Fig. 10c, f). Overall, PCOD-NN achieves agreement with the selected GIIRS observations comparable to RTTOV v14: its longwave RMSE is slightly larger, whereas its midwave RMSE is smaller. In both bands, the RMSE between the two simulations is smaller than either model's RMSE relative to GIIRS.

4.5 Cross-instrument application to FY-4B/AGRI

The instrument-independent optical-depth emulator was coupled to the spectral response functions of the seven FY-4B/AGRI infrared channels without retraining. To evaluate its performance against observations, we analysed an AGRI ocean scene

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

acquired on 29 April 2024 at 01:00–01:15 UTC (18–23°N, 124–130°E). We retained 9,011 pixels classified as certainly clear in three consecutive AGRI cloud-mask products and passing a 3 × 3-pixel neighbourhood screen. ERA5 atmospheric profiles and skin temperatures were matched to individual pixels. The simulations used pixel-dependent skin temperatures and a fixed surface emissivity of 0.97. Figure 11 shows three channels representing water-vapour (CH10, 6.9 μm), window (CH12, 8.5 μm), and CO_2 (CH15, 13.2 μm) sensitivity, together with the PCOD-NN and RTTOV v14 differences from AGRI observations. For CH10, the PCOD-NN bias and RMSE relative to AGRI were −0.03 and 0.94 K, compared with −0.27 and 0.98 K for RTTOV v14. The corresponding values were +0.66 and 0.70 K versus +0.57 and 0.62 K for CH12, and +1.51 and 1.54 K versus +1.31 and 1.35 K for CH15. Both simulations reproduce the broad spatial brightness-temperature patterns, although both exhibit a warm bias in CH15. Across all seven AGRI infrared channels, the RMSE ranges from 0.70 to 1.54 K for PCOD-NN and from 0.42 to 1.35 K for RTTOV v14. RTTOV has the lower RMSE, whereas PCOD-NN performs slightly better in CH10. Thus, this single-scene test demonstrates that the same PCOD-NN optical-depth emulator can be applied to AGRI observations without retraining.

4.6 Jacobian accuracy validation

Data assimilation and physical retrieval require the forward operator to provide accurate Jacobians (Rodgers, 2000; Bai et al., 2020). We obtain the Jacobian by automatic differentiation. The correctness of its software implementation has been confirmed by standard numerical checks; below we validate its physical accuracy against the LBLRTM finite-difference truth.

The BT-level Jacobian is the quantity actually used in assimilation (Yin et al., 2020). We take one clear-sky ERA5 profile. For each level, we perturb the temperature, the log specific humidity, or the trace-gas VMR and recompute the seven-gas mixed LBLRTM optical depth of the affected layers. This is passed through the same radiative transfer and instrument convolution to obtain the brightness temperature, and a central difference then gives the "truth" of each Jacobian. Figure 12 plots the Jacobian profiles for six representative variables, each in its diagnostic channel, in panel order: temperature $\partial BT/\partial T$ (a), CO_2 $\partial BT/\partial \ln$ VMR (b), water vapor $\partial BT/\partial \ln q$ (c), and $\partial BT/\partial \ln$ VMR for O_3 (d), CO (e), and N_2O (f). The automatic differentiation (red dashed line) and the LBLRTM finite difference (black solid line) almost completely overlap, with a median relative error at the high-sensitivity points of 0.1–1.6% in each panel (about 0.6% for temperature and about 0.4% for water vapor).

To confirm that this agreement holds across the atmospheric climatology rather than at a single state, we repeat the line-by-line finite-difference validation for 12 diverse profiles drawn from the SeeBor V5.1 representative set. These span tropical-to-polar conditions (latitude −69° to +76°, total precipitable water 0.07–4.5 cm, skin temperature 255–300 K, each profile with its own surface temperature), over all about 1682 GIIRS channels. Central finite differences use temperature perturbations of ±0.5 K and water-vapour VMR factors of $e^{\pm 0.02}$, and the relative-error statistics in Fig. 13 include only points whose absolute reference Jacobians exceed 10% of the respective global maximum for temperature and water vapour. Pooling every profile, channel, and level (Fig. 13), the automatic-differentiation and LBLRTM finite-difference Jacobians agree with $R^2 > 0.999$ and a regression slope of 1.00, with a mean bias of order 10^{-5}. The median relative error is about 0.2% for $\partial BT/\partial T$ and 0.3% for $\partial BT/\partial \ln q$, and the 95th-percentile relative error is 0.8% and 1.9%, respectively. The RMS error (1.3×10^{-3} K K^{-1} and 9.9×10^{-3}

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

K) is two to three orders of magnitude below the peak Jacobian, widening only modestly in the upper troposphere where the absorption lines are narrowest.

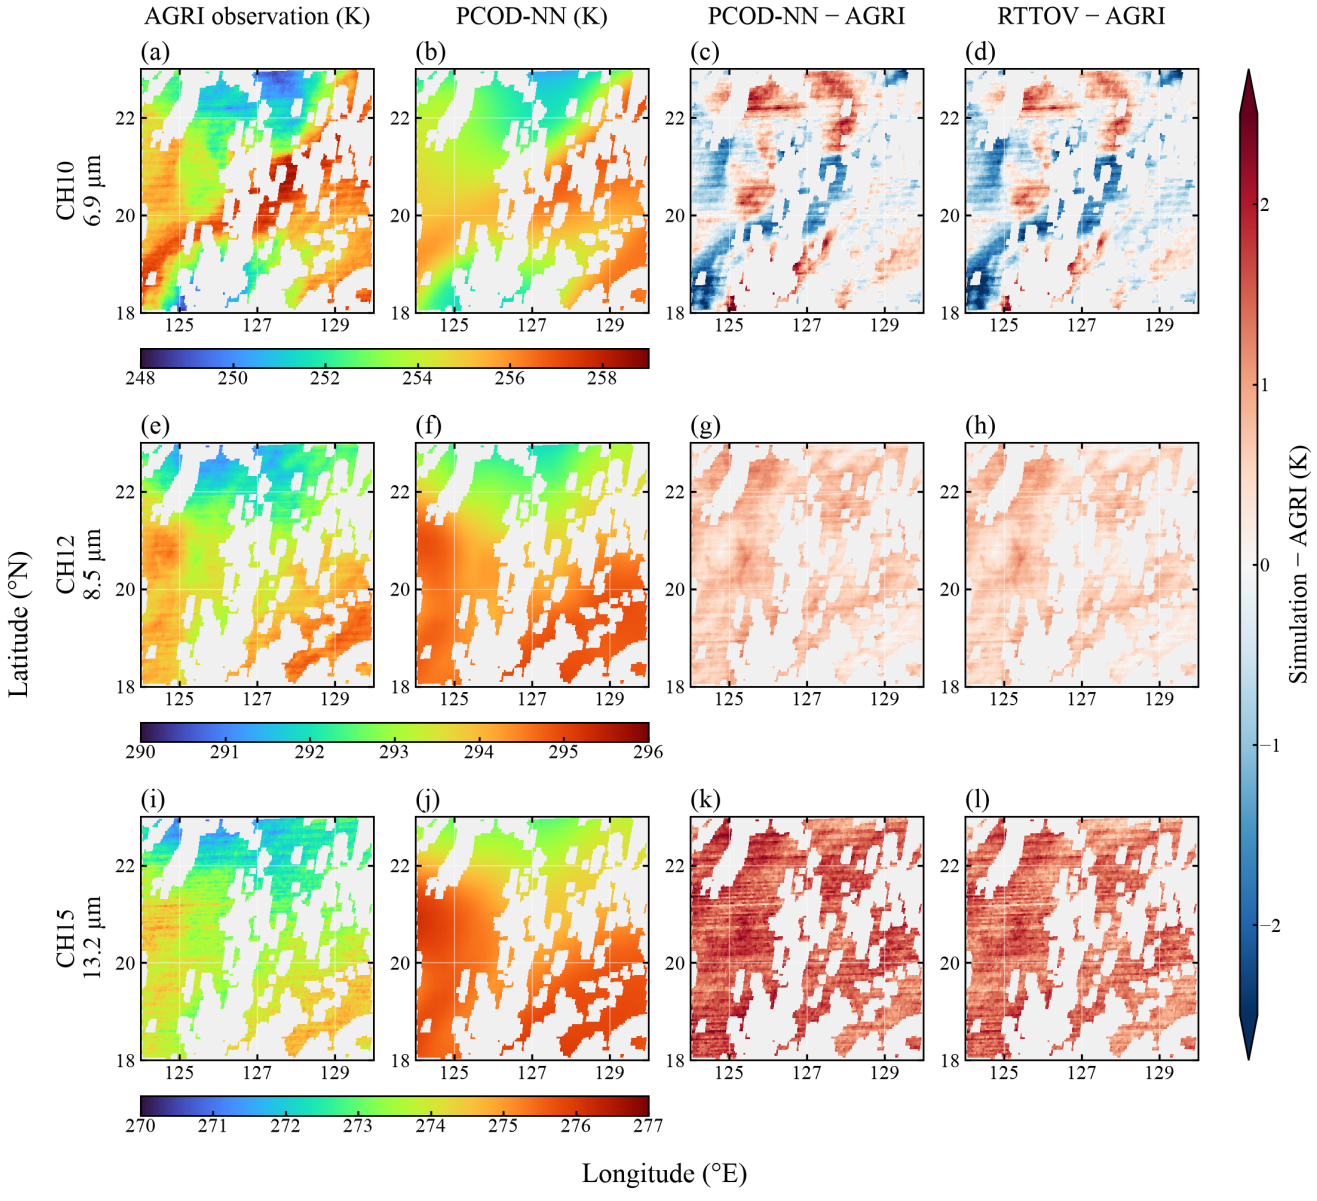


Figure 11. Spatial comparison of FY-4B/AGRI brightness temperatures over the ocean (18–23°N, 124–130°E) on 29 April 2024, 01:00–01:15 UTC. Rows show CH10 (6.9 µm; a–d), CH12 (8.5 µm; e–h), and CH15 (13.2 µm; i–l). Columns show the AGRI Level-1 observation, PCOD-NN simulation, PCOD-NN − AGRI difference, and RTTOV v14 − AGRI difference, respectively. The 9,011 displayed pixels passed the certain-clear criterion in three consecutive AGRI cloud-mask products and a 3 × 3-pixel neighborhood screen; blank areas were excluded. Observation and simulation share a brightness-temperature scale within each row, while all difference maps use the same ±2.5 K scale.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

By treating trace-gas concentrations as explicit inputs, PCOD-NN enables automatic differentiation of channel brightness temperatures with respect to ln VMR for CO_2, O_3, CO, and N_2O, as shown in Fig. 12. Emulators that use only temperature and water vapor as atmospheric inputs do not explicitly represent sensitivity to variations in these trace-gas concentrations. The resulting Jacobians provide sensitivity information for potential applications in trace-gas retrieval.

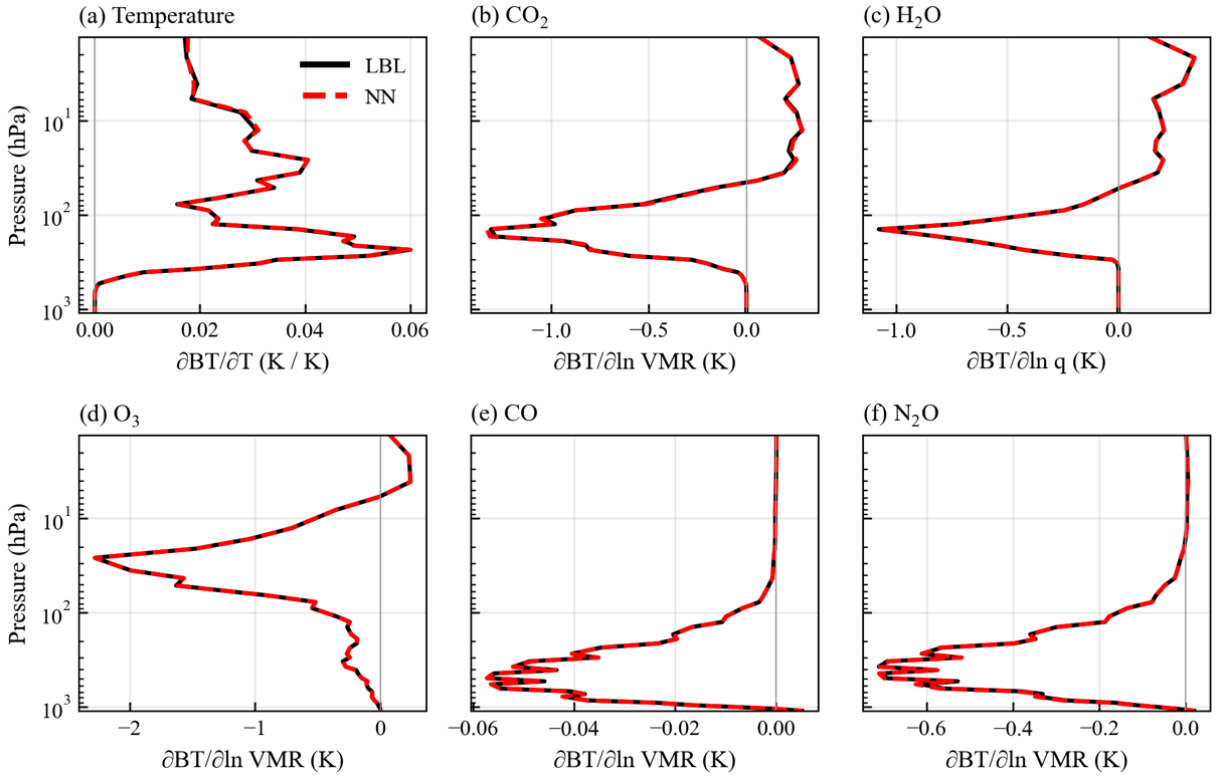


Figure 12. Representative BT Jacobians: automatic differentiation versus LBLRTM finite difference (single reference profile). The six panels (a)–(f) are the BT Jacobians of temperature, CO_2, water vapor, O_3, CO, and N_2O in their diagnostic channels (central wavenumbers 700, 700, 1700, 1040, 2150, 2200 cm^{-1}, respectively); (a) is ∂BT/∂T, (c) is ∂BT/∂ln q, and (b)(d)(e)(f) are ∂BT/∂ln VMR. The red dashed line is automatic differentiation (legend NN), the black solid line is LBLRTM finite difference (legend LBL), and the vertical axis is log pressure. The median relative error |Δ| at the high-sensitivity points of panels (a)–(f) is 0.6%, 1.6%, 0.4%, 0.1%, 0.6%, and 0.5%, respectively.

5 Discussion and conclusion

We have introduced and validated PCOD-NN, a neural-network emulator of layer-by-layer, gas-by-gas monochromatic optical depth over 600–2600 cm^{-1} at 0.01 cm^{-1} resolution. It predicts PCA-compressed ln(OD) spectra directly from the atmospheric state, without line-by-line absorption calculations during inference. On the held-out test sets, the mean relative OD error was below 0.2% for all seven gases. Because the emulated OD is instrument-independent, the same trained model ensemble can be coupled to different instruments without retraining the OD emulator, provided that their spectral responses lie within the

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

covered spectral range. This cross-instrument flexibility is demonstrated with FY-4B/GIIRS and FY-4B/AGRI. For 426 quality-screened clear-sky GIIRS fields of view, PCOD-NN achieved agreement with observations comparable to RTTOV v14. For 9,011 cloud-mask-selected clear-sky ocean pixels from one AGRI scene, its RMSE against Level-1 observations ranged from 0.70 to 1.54 K. The GPU-based OD computation was approximately 620–970 times faster than single-core LBLRTM. The differentiable pipeline also produced accurate temperature and water-vapor Jacobians across 12 diverse profiles, with median relative errors of 0.2% and 0.3%, respectively.

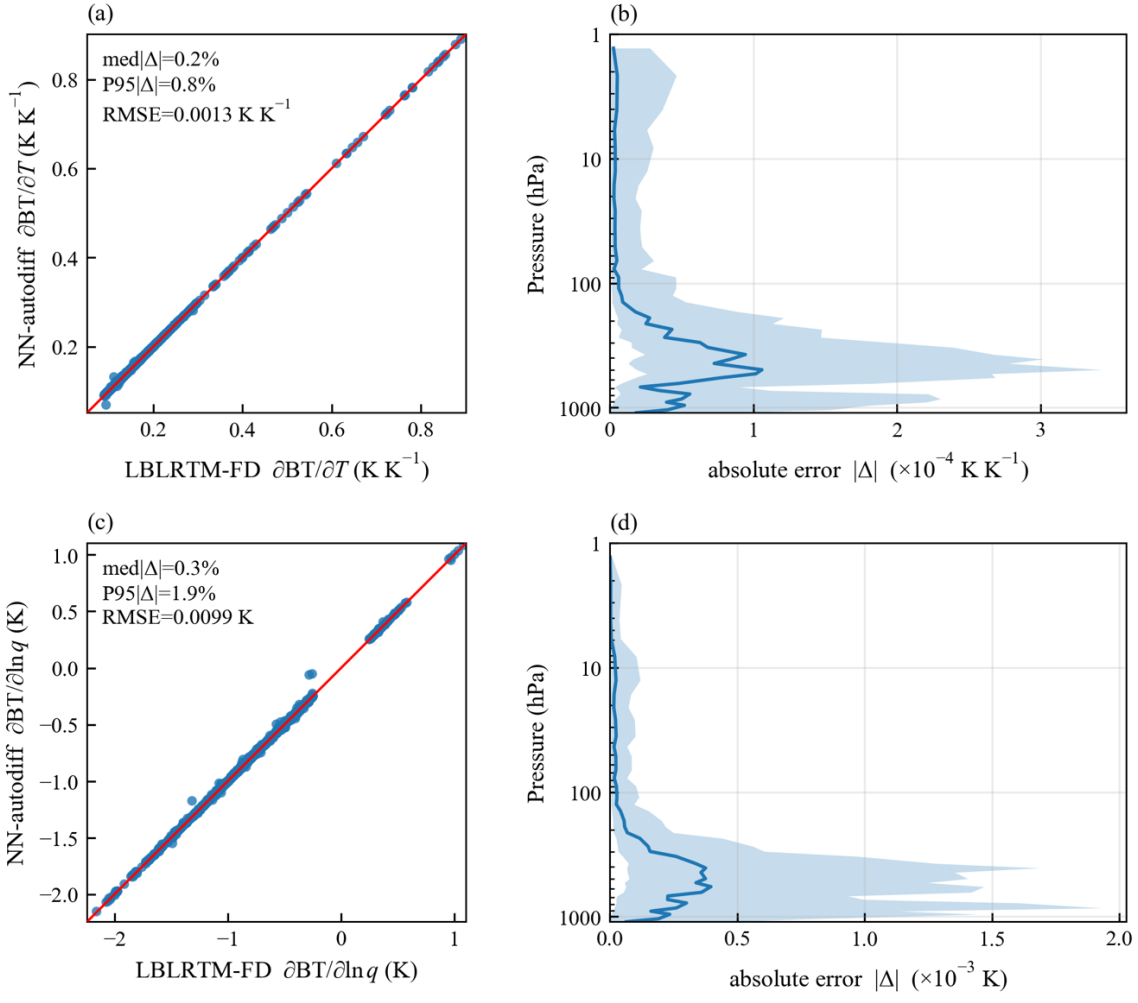


Figure 13. Brightness-temperature Jacobian validation across 12 diverse SeeBor profiles (tropical-to-polar; all about 1682 GIIRS channels): NN automatic differentiation versus LBLRTM finite difference. Top row, temperature Jacobian ∂BT/∂T; bottom row, water-vapor Jacobian ∂BT/∂ln q. (a, c) Pooled density scatter over all profiles, channels, and levels with the 1:1 line; the inset gives the median and 95th-percentile relative error and the RMS error. (b, d) Absolute error |Δ| versus pressure (median line and interquartile band). Pooled over all profiles, the automatic-differentiation and LBLRTM Jacobians agree with $R^2 > 0.999$, slope 1.00, median relative error 0.2% (∂BT/∂T) and 0.3% (∂BT/∂ln q), and 95th-percentile relative error 0.8% and 1.9%.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

Using OD as the emulation target creates a clear trade-off. It retains high-resolution spectral information and allows flexible coupling with different radiative-transfer solvers and instrument-response functions. However, the high-dimensional output requires PCA compression, substantial memory, and an external radiative-transfer and convolution procedure. For GIIRS channel simulation, the GPU implementation achieved end-to-end performance comparable to single-core RTTOV, although multithreaded RTTOV remained substantially faster. This comparison reflects different computational workloads. RTTOV predicts instrument channels directly, whereas PCOD-NN first calculates high-resolution monochromatic spectra. PCOD-NN is therefore complementary to channel-based models such as RTTOV and CRTM: it provides high-resolution, gas-resolved monochromatic optical depth and automatic-differentiation Jacobians, while allowing the trained OD emulator to be reused for a new instrument without retraining or generating instrument-specific transmittance coefficient files, provided that the instrument's ILS or SRF is supplied and lies within the model's spectral range.

Several limitations remain. The current model is restricted to clear-sky gaseous absorption over 600–2600 cm^{-1} and 1–1100 hPa. It does not include clouds, aerosols, or scattering and is not suitable for limb or middle-atmosphere applications. The model is also tied to the LBLRTM spectroscopic database and continuum configuration used for training, and its reported accuracy applies only within the sampled atmospheric domain. Nevertheless, the modular gas- and pressure-specific design facilitates extension to additional gases, pressure ranges, and spectral regions. Such extensions require new LBLRTM reference calculations, training, and independent validation. Overall, PCOD-NN provides a fast, flexible, and differentiable approximation of monochromatic optical depth. It offers a promising basis for high-resolution forward simulation and for future retrieval, sensitivity-analysis, and data-assimilation applications.

Code and data availability

The PCOD-NN model code (data generation, PCA fitting, training, inference, radiative transfer, instrument convolution, the automatic-differentiation Jacobian, and all figure-generation scripts), the 14 trained sub-models (seven gases × two pressure ranges, comprising the network weights and the PCA bases), the cached data underlying every figure, and a 50-sample-per-gas subset of the training database are archived at Zenodo (Yan et al., 2026; https://doi.org/10.5281/zenodo.20958781), together with the conda environment specification. The PCOD-NN code is released under the MIT licence, and the accompanying data (the trained sub-models, the cached figure data, and the training-data subset) under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence. The full training database (~430 GB of single-layer monochromatic optical-depth spectra at 0.01 cm^{-1} resolution) is too large to archive but is fully regenerable with the provided data-generation scripts driving LBLRTM. The forward model relies on third-party resources that are not redistributed here and retain their own licenses: LBLRTM v12.17 and the MT_CKD continuum model v4.3 (Clough et al., 2005; Mlawer et al., 2012; https://github.com/AER-RC/LBLRTM); ERA5 reanalysis profiles on pressure levels (Hersbach et al., 2020; Copernicus Climate Change Service Climate Data Store, https://doi.org/10.24381/cds.bd0915c6); SeeBor V5.1 profiles (CIMSS,

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

https://cimss.ssec.wisc.edu/training_data/); and FY-4B GIIRS/AGRI Level-1 data (National Satellite Meteorological Center, CMA, available on request from https://satellite.nsmc.org.cn).

Author contributions

J. Y.: Methodology, Software, Investigation, Validation, Visualization, Writing – original draft. C. L.: Conceptualization, Supervision, Funding acquisition, Writing – review & editing. Y. S.: Data curation, Investigation. D. D.: Resources, Formal analysis. P. Z.: Resources, Supervision.

Competing interests

The authors declare that they have no competing interests.

Acknowledgements

We gratefully acknowledge the providers of the third-party resources used in this study: AER for LBLRTM and the MT_CKD continuum model, the Copernicus Climate Data Store for the ERA5 reanalysis, CIMSS/SSEC for the SeeBor profile database, and the National Satellite Meteorological Center (CMA) for the FY-4B GIIRS/AGRI Level-1 data. The model simulations were conducted at the High Performance Computing Center of Nanjing University of Information Science & Technology.

Financial support

This research has been supported by the National Natural Science Foundation of China (grant no. 42375128) and the Postgraduate Research and Practice Innovation Program of Jiangsu Province (grant no. 26CXJH2401).

References

Alvarado, M. J., Payne, V. H., Mlawer, E. J., Uymin, G., Shephard, M. W., Cady-Pereira, K. E., Delamere, J. S., and Moncet, J.-L.: Performance of the Line-By-Line Radiative Transfer Model (LBLRTM) for temperature, water vapor, and trace gas retrievals: recent updates evaluated with IASI case studies, Atmos. Chem. Phys., 13, 6687–6711, https://doi.org/10.5194/acp-13-6687-2013, 2013.

Bai, W., Zhang, P., Zhang, W., Li, J., Ma, G., Qi, C., and Liu, H.: Jacobian matrix for near-infrared remote sensing based on vector radiative transfer model, Sci. China Earth Sci., 63, 1353–1365, https://doi.org/10.1007/s11430-019-9586-7, 2020.

Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M.: Automatic differentiation in machine learning: a survey, J. Mach. Learn. Res., 18, 1–43, https://doi.org/10.48550/arXiv.1502.05767, 2018.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

Clough, S. A., Shephard, M. W., Mlawer, E. J., Delamere, J. S., Iacono, M. J., Cady-Pereira, K., Boukabara, S., and Brown, P. D.: Atmospheric radiative transfer modeling: a summary of the AER codes, J. Quant. Spectrosc. Radiat. Transfer, 91, 233–244, https://doi.org/10.1016/j.jqsrt.2004.05.058, 2005.

Ding, S., Yang, P., Weng, F., Liu, Q., Han, Y., van Delst, P., Li, J., and Baum, B.: Validation of the community radiative transfer model, J. Quant. Spectrosc. Radiat. Transfer, 112, 1050–1064, https://doi.org/10.1016/j.jqsrt.2010.11.009, 2011.

Gordon, I. E., Rothman, L. S., Hargreaves, R. J., Hashemi, R., Karlovets, E. V., Skinner, F. M., Conway, E. K., Hill, C., Kochanov, R. V., Tan, Y., Wcislo, P., Finenko, A. A., Nelson, K., Bernath, P. F., Birk, M., Boudon, V., Campargue, A., Chance, K. V., Coustenis, A., Drouin, B. J., Flaud, J.-M., Gamache, R. R., Hodges, J. T., Jacquemart, D., Mlawer, E. J., Nikitin, A. V., Perevalov, V. I., Rotger, M., Tennyson, J., Toon, G. C., Tran, H., Tyuterev, V. G., Adkins, E. M., Baker, A., Barbe, A., Cane, E., Csaszar, A. G., Dudaryonok, A., Egorov, O., Fleisher, A. J., Fleurbaey, H., Foltynowicz, A., Furtenbacher, T., Harrison, J. J., Hartmann, J.-M., Horneman, V.-M., Huang, X., Karman, T., Karns, J., Kassi, S., Kleiner, I., Kofman, V., Kwabia-Tchana, F., Lavrentieva, N. N., Lee, T. J., Long, D. A., Lukashevskaya, A. A., Lyulin, O. M., Makhnev, V. Yu., Matt, W., Massie, S. T., Melosso, M., Mikhailenko, S. N., Mondelain, D., Muller, H. S. P., Naumenko, O. V., Perrin, A., Polyansky, O. L., Raddaoui, E., Raston, P. L., Reed, Z. D., Rey, M., Richard, C., Tobias, R., Sadiek, I., Schwenke, D. W., Starikova, E., Sung, K., Tamassia, F., Tashkun, S. A., Vander Auwera, J., Vasilenko, I. A., Vigasin, A. A., Villanueva, G. L., Vispoel, B., Wagner, G., Yachmenev, A., and Yurchenko, S. N.: The HITRAN2020 molecular spectroscopic database, J. Quant. Spectrosc. Radiat. Transfer, 277, 107949, https://doi.org/10.1016/j.jqsrt.2021.107949, 2022.

Hatfield, S., Chantry, M., Dueben, P., Lopez, P., Geer, A., and Palmer, T.: Building tangent-linear and adjoint models for data assimilation with neural networks, J. Adv. Model. Earth Syst., 13, e2021MS002521, https://doi.org/10.1029/2021MS002521, 2021.

Havemann, S., Thelen, J.-C., Taylor, J. P., and Harlow, R. C.: The Havemann-Taylor Fast Radiative Transfer Code (HT-FRTC): a multipurpose code based on principal components, J. Quant. Spectrosc. Radiat. Transfer, 220, 180–192, https://doi.org/10.1016/j.jqsrt.2018.09.008, 2018.

He, K., Zhang, X., Ren, S., and Sun, J.: Deep residual learning for image recognition, in: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, 770–778, https://doi.org/10.1109/CVPR.2016.90, 2016.

Hendrycks, D. and Gimpel, K.: Gaussian error linear units (GELUs), arXiv [preprint], https://doi.org/10.48550/arXiv.1606.08415, 2016.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horanyi, A., Munoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Holm, E., Janiskova, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thepaut, J.-N.: The ERA5 global reanalysis, Q. J. R. Meteorol. Soc., 146, 1999–2049, https://doi.org/10.1002/qj.3803, 2020.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.

Lacis, A. A. and Oinas, V.: A description of the correlated k distribution method for modeling nongray gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres, J. Geophys. Res., 96, 9027–9063, https://doi.org/10.1029/90JD01945, 1991.

Li, Y., Han, W., Duan, W., Li, Z., and Li, H.: A machine learning-based observation operator for FY-4B GIIRS brightness temperatures considering the uncertainty of label data, J. Geophys. Res. Mach. Learn. Comput., 2, e2024JH000449, https://doi.org/10.1029/2024JH000449, 2025.

Liang, H., Zhang, P., Chen, L., Ma, G., Bai, W., and Qi, C.: A Gradient Boosting Tree method for rapid calculation of level-to-space transmittances, Acta Meteorol. Sin., 78, 853–863, https://doi.org/10.11676/qxxb2020.055, 2020.

Liu, C., Yao, B., Natraj, V., Weng, F., Le, T., Shia, R.-L., and Yung, Y. L.: A spectral data compression (SDCOMP) radiative transfer model for high-spectral-resolution radiation simulations, J. Atmos. Sci., 77, 2055–2066, https://doi.org/10.1175/JAS-D-19-0238.1, 2020.

Liu, X., Smith, W. L., Zhou, D. K., and Larar, A.: Principal component-based radiative transfer model for hyperspectral sensors: theoretical concept, Appl. Opt., 45, 201–209, https://doi.org/10.1364/AO.45.000201, 2006.

Loshchilov, I. and Hutter, F.: Decoupled weight decay regularization, in: Int. Conf. Learn. Represent. (ICLR), New Orleans, LA, USA, https://doi.org/10.48550/arXiv.1711.05101, 2019.

Matricardi, M.: A principal component based version of the RTTOV fast radiative transfer model, Q. J. R. Meteorol. Soc., 136, 1823–1835, https://doi.org/10.1002/qj.680, 2010.

McKay, M. D., Beckman, R. J., and Conover, W. J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, Technometrics, 21, 239–245, https://doi.org/10.1080/00401706.1979.10489755, 1979.

Menzel, W. P., Schmit, T. J., Zhang, P., and Li, J.: Satellite-based atmospheric infrared sounder development and applications, Bull. Am. Meteorol. Soc., 99, 583–603, https://doi.org/10.1175/BAMS-D-16-0293.1, 2018.

Mlawer, E. J., Payne, V. H., Moncet, J.-L., Delamere, J. S., Alvarado, M. J., and Tobin, D. C.: Development and recent evaluation of the MT_CKD model of continuum absorption, Philos. Trans. R. Soc. A, 370, 2520–2556, https://doi.org/10.1098/rsta.2011.0295, 2012.

Natraj, V. and Spurr, R. J. D.: Accelerating radiative transfer calculations in the thermal infrared through principal component analysis of inherent optical properties, J. Quant. Spectrosc. Radiat. Transfer, 339, 109415, https://doi.org/10.1016/j.jqsrt.2025.109415, 2025.

Natraj, V., Jiang, X., Shia, R.-L., Huang, X., Margolis, J. S., and Yung, Y. L.: Application of principal component analysis to high spectral resolution radiative transfer: a case study of the O_2 A band, J. Quant. Spectrosc. Radiat. Transfer, 95, 539–556, https://doi.org/10.1016/j.jqsrt.2004.12.024, 2005.

Natraj, V., Shia, R.-L., and Yung, Y. L.: On the use of principal component analysis to speed up radiative transfer calculations, J. Quant. Spectrosc. Radiat. Transfer, 111, 810–816, https://doi.org/10.1016/j.jqsrt.2009.11.004, 2010.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.



Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.: PyTorch: an imperative style, high-performance deep learning library, in: Adv. Neural Inf. Process. Syst. (NeurIPS) 32, Vancouver, Canada, 8024–8035, https://doi.org/10.48550/arXiv.1912.01703, 2019.

Rodgers, C. D.: Inverse Methods for Atmospheric Sounding: Theory and Practice, Series on Atmospheric, Oceanic and Planetary Physics, Vol. 2, World Scientific, Singapore, https://doi.org/10.1142/3171, 2000.

Saunders, R., Hocking, J., Turner, E., Rayer, P., Rundle, D., Brunel, P., Vidot, J., Roquet, P., Matricardi, M., Geer, A., Bormann, N., and Lupu, C.: An update on the RTTOV fast radiative transfer model (currently at version 12), Geosci. Model Dev., 11, 2717–2737, https://doi.org/10.5194/gmd-11-2717-2018, 2018.

Smith, L. N. and Topin, N.: Super-convergence: very fast training of neural networks using large learning rates, in: Proc. SPIE 11006, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications, 1100612, https://doi.org/10.1117/12.2520589, 2019.

Spurr, R., Natraj, V., Lerot, C., Van Roozendael, M., and Loyola, D.: Linearization of the principal component analysis method for radiative transfer acceleration: application to retrieval algorithms and sensitivity studies, J. Quant. Spectrosc. Radiat. Transfer, 125, 1–17, https://doi.org/10.1016/j.jqsrt.2013.04.002, 2013.

Stegmann, P. G., Johnson, B., Moradi, I., Karpowicz, B., and McCarty, W.: A deep learning approach to fast radiative transfer, J. Quant. Spectrosc. Radiat. Transfer, 280, 108088, https://doi.org/10.1016/j.jqsrt.2022.108088, 2022.

Turner, D. D., Tobin, D. C., Clough, S. A., Brown, P. D., Ellingson, R. G., Mlawer, E. J., Knuteson, R. O., Revercomb, H. E., Shippert, T. R., Smith, W. L., and Shephard, M. W.: The QME AERI LBLRTM: a closure experiment for downwelling high spectral resolution infrared radiance, J. Atmos. Sci., 61, 2657–2675, https://doi.org/10.1175/JAS3300.1, 2004.

Ukkonen, P.: Exploring pathways to more accurate machine learning emulation of atmospheric radiative transfer, J. Adv. Model. Earth Syst., 14, e2021MS002875, https://doi.org/10.1029/2021MS002875, 2022.

Ukkonen, P., Pincus, R., Hogan, R. J., Nielsen, K. P., and Kaas, E.: Accelerating radiation computations for dynamical models with targeted machine learning and code optimization, J. Adv. Model. Earth Syst., 12, e2020MS002226, https://doi.org/10.1029/2020MS002226, 2020.

Veerman, M. A., Pincus, R., Stoffer, R., van Leeuwen, C. M., Podareanu, D., and van Heerwaarden, C. C.: Predicting atmospheric optical properties for radiative transfer computations using neural networks, Philos. Trans. R. Soc. A, 379, 20200095, https://doi.org/10.1098/rsta.2020.0095, 2021.

Yan, J., Liu, C., Sun, Y., Di, D., and Zhang, P.: PCOD-NN v1.0: a differentiable optical-depth emulator for line-by-line infrared radiative transfer, Zenodo [code and data set], https://doi.org/10.5281/zenodo.20958781, 2026.

Yang, J., Zhang, Z., Wei, C., Lu, F., and Guo, Q.: Introducing the new generation of Chinese geostationary weather satellites, Fengyun-4, Bull. Am. Meteorol. Soc., 98, 1637–1658, https://doi.org/10.1175/BAMS-D-16-0065.1, 2017.

https://doi.org/10.5194/egusphere-2026-6064
Preprint. Discussion started: 13 November 2026
© Author(s) 2026. CC BY 4.0 License.



Yin, R., Han, W., Gao, Z., and Di, D.: The evaluation of FY-4A's Geostationary Interferometric Infrared Sounder (GIIRS) long-wave temperature sounding channels using the GRAPES global 4D-Var, Q. J. R. Meteorol. Soc., 146, 1459–1476, https://doi.org/10.1002/qj.3746, 2020.