



A machine-learning-based marine atmosphere boundary layer (MABL) moisture profile retrieval product from GNSS-RO deep refraction signals

Jie Gong¹, Dong L. Wu¹, Michelle Badalov², Manisha Ganeshan^{3,1}, and Minghua Zheng⁴

¹Climate and Radiation Lab, NASA Goddard Space Flight Center, Greenbelt, MD, USA

²Dept. of Computer Science, Univ. of Maryland, College Park, MD, USA

³GESTAR-II/Morgan State University, Baltimore, MD, USA

⁴Scripps Institution of Oceanography, University of California at San Diego, La Jolla, CA, USA

Correspondence: Jie Gong (jie.gong@nasa.gov)

Received: 31 March 2024 – Discussion started: 20 June 2024

Revised: 3 March 2025 – Accepted: 13 June 2025 – Published: 27 August 2025

Abstract. Marine atmosphere boundary layer (MABL) water vapor amount and gradient impact global energy transport through directly affecting the sensible and latent heat exchange between the ocean and atmosphere. Yet, it is a well-known challenge for satellite remote sensing to profile MABL water vapor, especially when cloud or a sharp vertical gradient of water vapor is present. Wu et al. (2022) identified good correlations between the Global Navigation Satellite System (GNSS) deep refraction signal-to-noise-ratio (SNR) value and the global MABL water vapor specific humidity when the radio occultation (RO) signal is ducted by the moist planetary boundary layer (PBL), and they laid out the underlying physical mechanisms to explain such a correlation. In this work, we apply a machine learning/artificial intelligence (ML/AI) technique to demonstrate the feasibility of profile-by-profile MABL water vapor retrieval using the SNR signal. Three convolutional neural network (CNN) models are trained using multi-months of global collocated hourly ERA-5 reanalysis and COSMIC-1, Metop-A, and Metop-B 1 Hz SNR observations between 975–850 hPa with 25 hPa vertical resolution. The COSMIC-1 ML model is then applied to both COSMIC-1 and COSMIC-2 in other time ranges for independent retrieval and validation. The Monte Carlo Dropout method was employed for the uncertainty estimation. Comparison against multiple field campaign radiosonde/dropsonde observations globally suggests that SNR-ML-method-retrieved water vapor consistently outperforms the wetPrf/wetPf2 standard retrieval product at all six pressure

levels between 975 and 850 hPa and either outperforms or achieves similar performance against ERA-5, indicating real and useful information is gained from the SNR signal, though training was performed against the reanalysis. The climatology and diurnal cycle of MABL structure constructed from the SNR-ML technique are studied and compared to the reanalysis. Disparities of climatology suggest ERA-5 may systematically produce dry biases at high latitudes and wet biases in marine stratocumulus regions. The diurnal cycle amplitudes are too weak and sometimes off phase in ERA-5, especially in the Arctic and stratocumulus regions. These areas are particularly prone to PBL processes, where this GNSS SNR-ML water vapor product may contribute the most.

1 Introduction

As a key component of the Earth's lower atmosphere, planetary boundary layer (PBL) water vapor plays a pivotal role in the Earth's energy budget, exerting a profound influence on weather and climate processes. It is an essential factor of the Earth's energy budget, influencing radiative forcing and consequently climate variability and long-term changes. Furthermore, PBL water vapor is instrumental in modulating local and regional weather patterns by affecting cloud formation, precipitation, and temperature. Therefore, the study of PBL water vapor stands as a vital element in advancing

our comprehension of the Earth's atmosphere and its broader implications for our planet's climate system.

70 % of the Earth's surface is covered by water. The sensible and latent heat exchange between the ocean boundary and the marine atmosphere boundary layer (MABL) happens at different spatial and temporal scales, which is determined not only by ocean surface properties (e.g., wind speed, sea surface temperature) but also by MABL thermodynamic structures. For example, in the context of susceptibility of polar area to the climate change, Boisvert et al. (2015) found that Arctic PBL humidity and temperature biases in the reanalysis are the major error sources for the evaporation estimation compared to satellite observations. Cloud–climate feedback is another motivation highlighted by NASA's PBL incubation study (Teixeira et al., 2021). As another example, Millán et al. (2019) found strong correlation between MABL cloud top height and below-cloud water vapor amount using two joint satellite retrieval products.

Data sparsity is a critical problem for advancing MABL science. Satellite remote sensing undoubtedly provides the best solution in terms of global coverage, but it is very difficult to retrieve MABL water vapor (WV) and its vertical distribution when cloud or sea ice is present. When clouds are present in the scene, emissions from clouds often overwhelm the emission signal from the MABL water vapor and prevent passive instruments sensing the below-cloud atmosphere. When sea ice is present, scattering or surface emission from the sea ice is often inseparable from water vapor emission signals and distorts the retrieval result. Taking the aforementioned two research studies as examples, Boisvert et al. (2015) use Level-2 AIRS water vapor and temperature retrieval products, which are only available for clear or partially cloudy sky situations, so they inherently contain a sampling bias. Millán et al. (2019) derived MABL total WV amount from subtracting MODIS above-cloud water vapor from AMSU-A total column water vapor, which still lacks the vertical information of WV in the MABL.

Using the low-frequency microwave L-band to transmit signals along the limb path, the Global Navigation Satellite System (GNSS) satellite overcomes the two difficulties above and provides high vertical resolution (100–200 m) of the MABL water vapor under all-sky conditions. GNSS Radio Occultation (GNSS-RO) retrieves temperature and water vapor profiles using the 1D-Var approach routinely from the Level 2 bending angle product (referred as “standard L2 product” or “operational L2 product” hereafter), the latter of which is used operationally in numerical weather data assimilation systems to improve weather forecasts (e.g., Kuo et al., 2000). Because of the rapid growth of SmallSat/CubeSat constellations from both the commercial and the non-profit sectors, the GNSS-RO technique provides a promising future for the needed global spatial–temporal sampling of MABL WV and its variability. Like other limb sounders, the disadvantage of GNSS-RO is its relatively coarse horizontal resolution (several hundred kilometers) that smears out horizon-

tally inhomogeneous signals. This is typically not a big concern in MABL as the vertical gradient is much sharper than the horizontal gradient and harder to characterize.

However, GNSS-RO WV retrieval profiles have excessively high failure rate in the MABL. That is because the GNSS-RO signal-to-noise ratio (SNR) decreases with decreasing altitude due to the atmospheric defocusing effect, and the Level-2 RO signal hence often does not meet the SNR threshold near the surface. As a result, the GNSS-RO 1D-Var-based retrievals often fail in the MABL due to weak RO signals. Figure 1 gives an example of the success statistics (%) as a function of height for temperature (Fig. 1a) and water vapor (Fig. 1b) over the tropical ocean (10° S–10° N). Using 0.5 and 1 km above the ocean surface as the reference lines, we can see that although the COSMIC-2 (Constellation Observing System for Meteorology, Ionosphere, and Climate-2) has significantly improved its SNR compared to its predecessor COSMIC-1, the success rate is still about 60 % at 0.5 km and slightly over 70 % at 1 km for the GNSS-RO WV retrieval, while this number is only 40 % and 55 % for COSMIC-1 at respective altitudes. The low SNR widely exists for commercial GNSS satellites, especially in the lowest 500 m above the surface (Ganeshan et al., 2025). Moreover, even past the SNR threshold, some bending angle profiles are significantly biased in the PBL when ducting happens because the refractivity index becomes negative, which leads to biases in the operational water vapor retrievals (Feng et al., 2020).

Wu et al. (2022) found that the Level-1B deep SNR from the straight-line height (H_{SL}) is statistically significantly correlated with the MABL water vapor amount in the European Centre for Medium-Range Weather Forecasts (ECMWF) Reanalysis v5 (ERA-5) after averaging over a month at $2.5^\circ \times 2.5^\circ$ grid resolution. The averaging is necessary to effectively beat down the random noise. This paper attributed such a positive correlation to the strong refraction from a horizontally stratiform and dynamically quiet MABL water vapor layer that acts to enhance the SNR amplitude at deep H_{SL} through ducting and diffraction/interference (a summary recapitulation of the physical mechanism can be found in Sect. 2.3). Some caveats of this work limit its application to weather phenomena. First, it builds upon a single-level regression statistics, the correlation coefficient of which was found to be the highest at $H_{SL} = -100$ km in the tropics and $H_{SL} = -80$ km at high latitudes. Hence, any simple linear-regression-based retrieval algorithm will suffer from arbitrary latitudinal discontinuities. As a matter of fact, SNRs at different H_{SL} levels are found correlated with MABL water vapor with different signs and magnitudes (e.g., Fig. 2), which should be used together to enhance the information content. Secondly, the robust relationship is only found for monthly averages in Wu et al. (2022) because the profile-by-profile noise is usually too high to yield a meaningful retrieval from SNR, and only through averaging large amount of profiles can the noise be lowered down to the level where

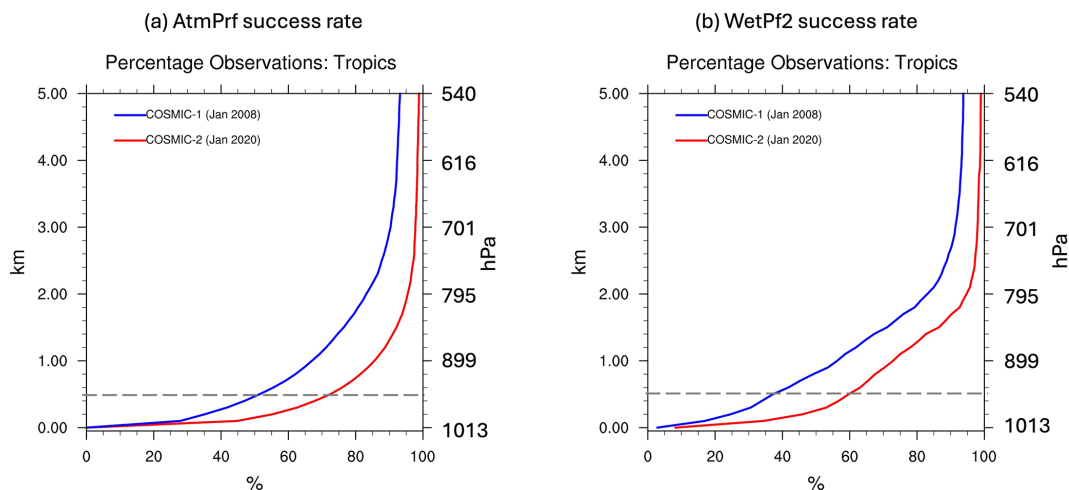


Figure 1. Level-2 atmPrf (temperature) and wetPf2 (water vapor) successful retrieval rate (%) as a function of height above sea level from COSMIC-1 during January 2008 (blue) and COSMIC-2 during January 2020 (red). The success rate is calculated by dividing the number of valid GNSS-RO retrieval files by the number of Level-1B files at a certain height. The dashed gray lines mark the reference at 0.5 km from the tropical ocean surface.

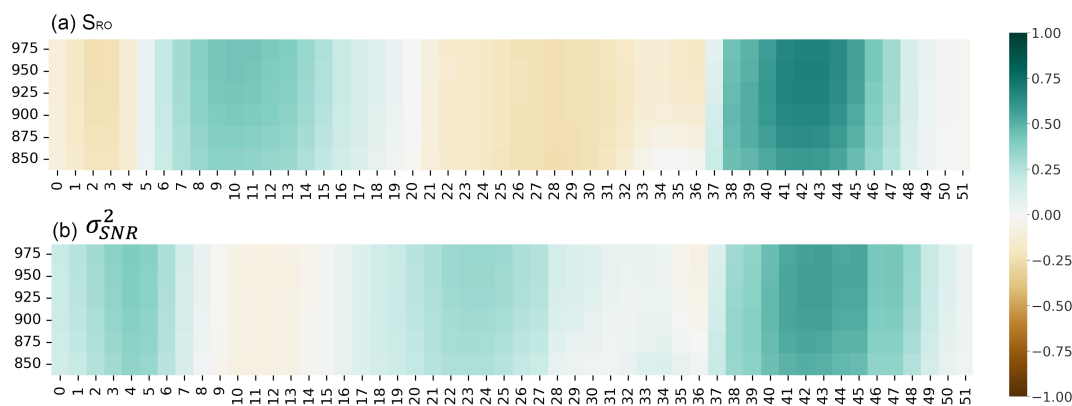


Figure 2. Correlation between collocated ERA-5 specific humidity at 975–850 hPa and S_{RO} (a) and σ_{SNR}^2 (b) at various excess phase levels from the training COSMIC-1-ERA-5 dataset). Only grid indices are shown in the axis titles, and the corresponding $\text{Log}_{10}(\phi_{L1})$ values can be found in Table A2.

the signal stands out. These are all caveats of a traditional statistical approach. The machine learning approach, however, is suitable at picking up weak signals through a large number of training data. As such, the scopes of this paper are to demonstrate the feasibility of using the ML method to extract MABL WV information from the GNSS SNR signals and to demonstrate the scientific value of this new product over the existing operational water vapor retrievals.

Artificial intelligence/machine learning (AI/ML) applications in remote sensing field have been trending in the last decade. They have been increasingly used in remote sensing fields in recent years. Traditional physics-based radiative transfer (RT) theories and models are used to link the remote sensing measurements (e.g., GNSS radio occultation signal) to the physical quantities (e.g., temperature and water vapor profiles). They are often highly non-linear and computation-

ally expensive and involve many explicit or embedded assumptions/simplifications, which may or may not propagate properly into part of the retrieval errors eventually. Given the fact that satellite measurements usually contain a large number of data, and the association is highly non-linear between the measurement space and the physical space, the retrieval process becomes an ideal test bed for ML capabilities. Some pioneer works have attempted this approach to retrieve PBL atmosphere profiles and achieved notable success. For example, Ye et al. (2021) used the routine radiosonde measurement at an Atmospheric Radiation Measurement (ARM) site as the ground truth to train a ground-based infrared spectrometer to predict the PBL height. The capability is limited to only the stations where both observations are routinely available. Milestein and Blackwell (2016) employed a neural network (NN) framework for retrieving the temperature and wa-

ter vapor profiles from the spaceborne Atmospheric Infrared Sounder (AIRS) observations (AIRS Version 7 product). The training “truth” was from the ECMWF analysis fields. It is worth mentioning that Milestein (2022), in a follow-up work, pointed out that the ML-only retrieval framework tends to smooth out sharp gradient features in proximity to the PBL top. To mitigate this caveat, Milstein et al. (2023) employ the 3D deep neural network training on the AIRS granule image against ERA-5 reanalysis that helps PBL height recognition from passive imagers.

In this paper, we will explore the ML capability at retrieving the MABL WV information from the deep SNR signal at profile-by-profile basis (i.e., Level-2 standard). Section 2 introduces the training and validation datasets as well as the model structure; Sect. 3 presents the retrieval results and independent validation; Sect. 4 expands the discussion to the usage of this product in studying MABL water vapor climatology and diurnal variabilities; and Sect. 5 summarizes the major findings and shortcomings of the current work that may be improved in the future.

2 Data and model

This section introduces the training, validation, and independent validation datasets, as well as the ML model architecture and the underlying physical foundations, that the ML technique is rooted in.

2.1 Training and validation datasets

The definition of SNR follows Wu et al. (2022), which uses the normalized SNR (S_{RO}):

$$S_{RO} = (SNR - \sigma) / (SNR_0 - \sigma) \quad (1)$$

$$\sigma_S^2 = \text{VAR}(S_{RO} - \overline{S_{RO}}). \quad (2)$$

SNR_0 is the free-atmosphere SNR. In practice, we use averaged SNR between the 35 and 65 km altitude range as the SNR_0 , and any profile with $SNR_0 < 200$ or $\sigma_{SNR_0}^2 > 0.05$ is considered to be of “low signal” and is filtered out. σ is the instrument-specific noise determined for each individual instrument from very deep H_{SL} . The value for σ used in this work is an updated version from Table A1 in Wu et al. (2022) and shown in Appendix A (Table A2). Wu et al. (2022) also found an instrument-dependent shift of the mean S_{RO} profile as a function H_{SL} . Luckily, such an issue can be resolved using the excess phase at L1 (ϕ_{L1}) as the vertical coordinate. In practice, the raw calculated S_{RO} and σ_S^2 are mapped to a fixed 52-level $\text{Log}_{10}(\phi_{L1})$ vertical grid. They are roughly linearly correlated with H_{SL} . The value for the vertical grid is listed in Table A2 in the Appendix A. In practice, we also filtered out bad open-loop profiles, profiles with data gap greater than 2 km, and profiles with outlier S_{RO} or σ_S^2 values.

The ERA-5 reanalysis is so far the best global reanalysis dataset in terms of PBL water vapor amount and dis-

tribution. Johnston et al. (2021) compared specific humidity from ERA-5 and MERRA-2 reanalysis against collocated and coincident GNSS-RO wetP2 specific humidity retrieval profiles and found ERA-5 outperforms MERRA-2 everywhere in the PBL. They both exhibit consistent dry biases with larger bias from mid-latitudes to high latitudes. However, ERA-5 percentage bias is roughly half of that from the MERRA-2 reanalysis in the PBL and tropopause layers. Given that many previous works used ERA-5 reanalysis or ECMWF analysis for training or validating the satellite retrievals for water vapor (e.g., Milestein and Blackwell, 2016; Milstein et al., 2023), especially some recent ones using them as the standard to evaluate recent GNSS-RO missions (e.g., Chang et al., 2022; Zhuran, 2023; Ganeshan et al., 2025), it is well justified to use ERA-5 hourly reanalysis as the “training” dataset to create a large sample globally. However, it is also warned in Johnston et al. (2021) that GNSS-RO retrievals tend to have their own biases especially in MABL, and in fact some other research suggested wet biases in certain regions (e.g., Virman et al., 2021).

In this work, we created a collocated and coincident ERA-5-SNR training and validation dataset. The SNR records are from four satellite series: COSMIC-1, COSMIC-2, Metop-A, and Metop-B. The periods for training, independent testing, and prediction are listed in Table 1. Note that the testing period is independent from the training period to avoid potential self-correlation using standard random splitting procedure. The prediction period however covers both training and validation periods, mainly for generating enough samples to construct statistically robust climatology (e.g., diurnal cycles). This however creates an unfortunate data leakage concern (e.g., as pointed out by Kapoor and Narayanan, 2023) for the comparison with the MAGIC campaign but not for the rest of other independent validation datasets (Table 2). The target variables are specific humidity at the aforementioned six pressure levels (975, 950, 925, 900, 875, and 850 hPa). The input parameters are 52 levels of S_{RO} , 52 levels of σ_S^2 , latitude, longitude, month, and rising/setting flag.

Figure 2 elucidates the linear correlation between COSMIC-1 S_{RO} at each of the 52 levels and ERA-5 specific humidity at 975, 950, 925, 900, 875, and 850 hPa over global ocean. The largest positive correlations are found around Level #40 to Level #45, which roughly correspond to $H_{SL} = -100$ to -80 km (Table A2). Based on the monthly averages, Wu et al. (2022) found the highest correlation at $H_{SL} = -100$ km in the tropics and at $H_{SL} = -80$ km for the polar regions, which is consistent with our profile-by-profile correlation as well. But Fig. 2 also shows positive or negative correlations at different $\text{Log}_{10}(\phi_{L1})$ levels, which prevent methods like multi-variable linear regression from working. σ_{SNR}^2 also exhibits non-linear patterns with slightly weaker correlations with MABL water vapor that are opposite in sign compared to those of S_{RO} . It is worth noting that these relationships are also instrument dependent, as can be clearly seen in the S_{RO} cross-correlation for Metop-

Table 1. Training, testing, and prediction periods. Periods are given in yyyy.mm–yyyy.mm.

Training (90 % and 10 % random-splitting)	COSMIC-1	2012.01–2012.12, 2016.01–2016.03, 2017.01–2017.03
	Metop-A	2017.01–2017.03
	Metop-B	2017.01–2017.03
Testing	COSMIC-1	2018.01–2018.03
	Metop-A	2018.01–2018.03
	Metop-B	2018.01–2018.03
Prediction	COSMIC-1	2012.01–2012.12, 2013.01–2013.12, 2016.01–2016.03, 2017.01–2017.03, 2018.01–2018.03
	COSMIC-2	2020.01–2020.12
	Metop-A	2012.01–2012.12, 2013.01–2013.12
	Metop-B	2013.02–2013.12

Table 2. Campaign information.

Campaign name	Period used for validation (yyyy.mm–yyyy.mm)	Location	Weather regime	Type	Reference
MARCUS	2017.11–2018.03	Southern Ocean	Mixed-phase PBL cloud	Radiosonde	Keeler et al. (2022)
ATOMIC	2020.01–2020.02	Tropical North Atlantic	Tropical trade wind zone	Radiosonde and Dropsonde	George et al. (2021)
EUREC4A	2020.01–2020.02	Tropical North Atlantic	Tropical trade wind zone	Radiosonde	Stephan et al. (2021)
MAGIC	2012.10–2013.09	Eastern North Pacific Ocean	Subtropical MABL	Radiosonde	Keeler et al. (2022)
ARRecon	2018.02; 2020.01–2020.02	Northeast Pacific off the coast of California	Atmospheric river	Dropsonde	Zheng et al. (2024)

A and Metop-B in the Appendix Figs. A1 and A2. Considering the instrument-dependent correlation patterns, three ML models are developed separately for COSMIC, Metop-A, and Metop-B satellites, although it is probably redundant to build two separate ML models for Metop-A and Metop-B separately as their correlation patterns are nearly identical. For the COSMIC series, we observed a similar pattern from COSMIC-2 compared to Fig. 2 after downsampling the frequency to 1 Hz (not shown). Therefore, the ML model developed using COSMIC-1 observations is applied directly to the downsampled COSMIC-2 SNR observations. Through this practice we can also test the transfer learning among similar satellite series for the hope of stitching them together for longer record in the future research.

The correlation holds with the same slope at a piece-wise level using individual profiles. For example, between SNR at $H_{SL} = -100$ km and ERA-5 specific humidity at 950 hPa,

Wu et al. (2022) observed the near linear correlation with monthly averaged and gridded data, while we can see that the same slope is preserved at a profile-by-profile level in Fig. 3. While this robust correlation proves that developing a Level-2 MABL specific humidity retrieval product using SNR profiles is feasible, the discernible larger noise at individual profile level versus month averages (Fig. 3d) suggests it is a challenging task. The ML method is hence introduced to tackle this highly complex regression problem.

GNSS-RO operational water vapor retrieval product provided by the University Corporation for Atmospheric Research (UCAR) is employed to evaluate the quality of the SNR-ML retrievals. This operational product is called “wetPf2”. Compared to an old processed “wetPrf” version from 2013, “wetPf2” has better penetration depth (Wee et al., 2022) and is used for constructing Fig. 1, but the “wet-Prf” product is used for the MAGIC campaign comparison

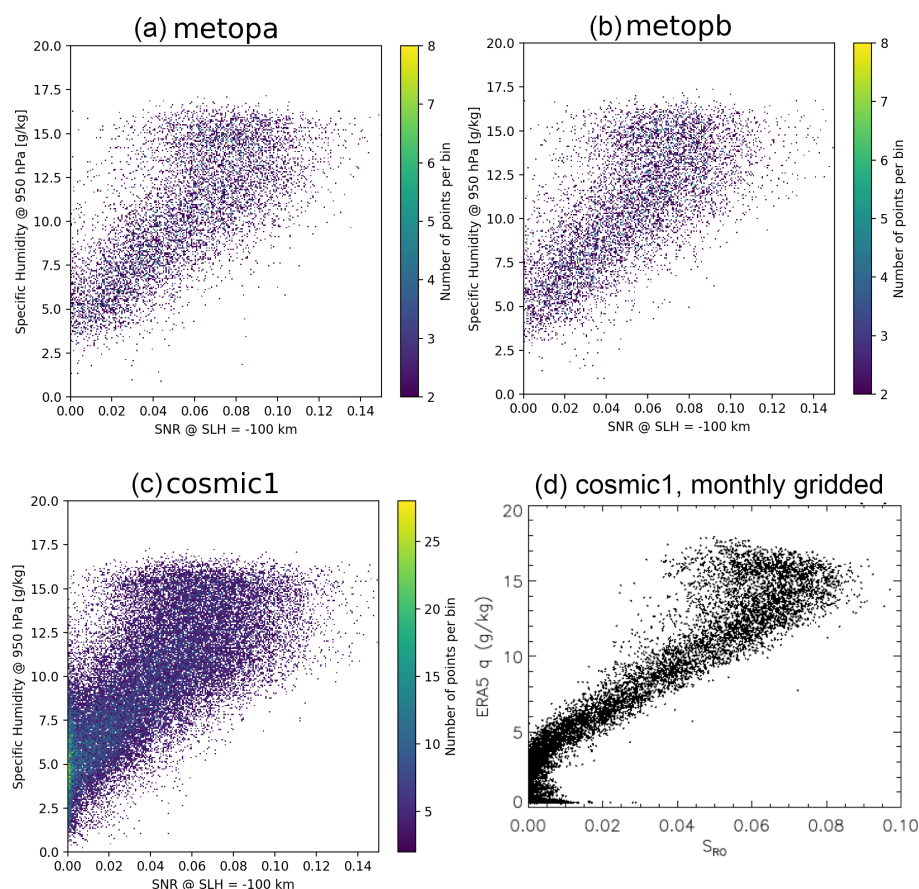


Figure 3. Density plots of the SNR-specific humidity relationship for (a) Metop-A, (b) Metop-B, and (c) COSMIC-1 constructed from the entire training dataset between 45° S and 45° N. The SNR value is taken from $H_{\text{SL}} = -100$ km, while the specific humidity value is taken at 950 hPa. Figure 9c from Wu et al. (2022) is reproduced here as panel (d) to demonstrate that the same relationship with the same slope holds at an individual profile level.

because of data availability constraints at the time when this research was conducted. We compared the success rate in the MABL between wetPrf and wetPf2 during January 2008 (not shown) and only found very marginal improvements for COSMIC-1. Note that the key Level-2 profile to enable the 1D-VAR retrieval used by the wetPrf/wetPf2 product is the bending angle, which is assimilated in the ERA-5 reanalysis. Therefore, this is not an independent evaluation dataset. The purpose of this comparison is to identify the merits and caveats of the SNR-ML retrievals against an existing mature product.

In addition to the independent testing, which is a standard procedure for ML/AI training and evaluation against the wetPrf/wetPf2 operational product, a handful of shipborne radiosonde campaign and airborne dropsonde campaign data are collected for further independent assessment. The campaign names, location, and total number of valid profiles are presented in Fig. 4 and Table 2. We can see from the summary of weather scenarios during each campaign that this independent validation dataset comprehensively covers ma-

jor marine weather regimes from the extremely dry Southern Ocean (MARCUS), mid-latitude stratocumulus region (MAGIC), and tropical trade cumulus region (EUREC4A, ATOMIC) to episodically wet atmospheric river events (AR-Recon). This exercise is critical for assessing the quality of ERA-5, the wetPrf/wetPf2 retrieval, and the Level 1 SNR-based retrieval under different weather scenarios. Moreover, as the ML model trained solely on COSMIC-1 SNR data is then applied to the COSMIC-2 data, the independent validation using the three campaigns in 2020 (ARRecon-2020, EUREC4A and ATOMIC) provides some solid evidence to evaluate the robustness of the “transfer learning”.

2.2 Machine learning model selection

The convolutional neural network (CNN) model (LeCun et al., 2015) is chosen as our regression ML model. The model internal architecture is illustrated in Fig. 5. There are a total of 109 input parameters, including one dimensional array of S_{RO} of 52 elements, one dimensional array of σ_S^2 of 52 elements, both interpolated to a fixed excess phase grid

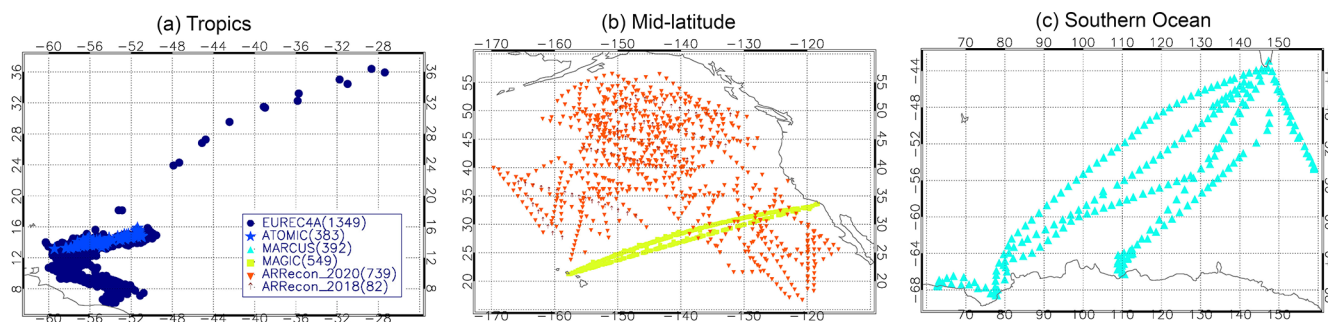


Figure 4. Maps for radiosonde/dropsonde locations from different shipborne or airborne campaigns in (a) the tropics, (b) the mid-latitudes, and (c) the Southern Ocean. Detailed campaign information can be found in Table 2. The total number of valid radiosonde/dropsonde profiles is listed in the parentheses in the legends.

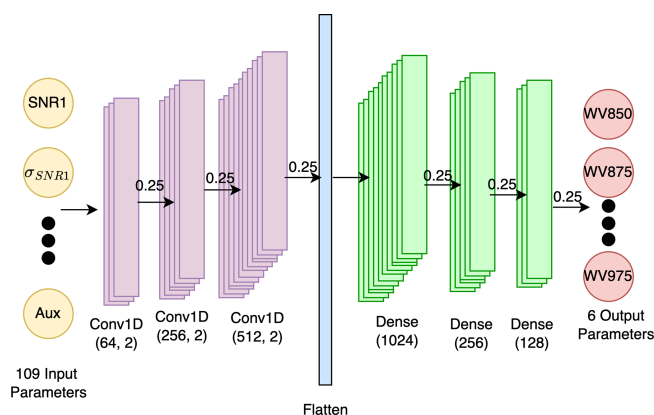


Figure 5. CNN model internal structure for this work. The numbers above the right-pointing arrows are the Monte Carlo dropout value applied between each layer. Numbers inside the parentheses of Conv1D layer indicate filter size and pool size, while numbers inside the parentheses of the dense layer indicate the number of fully connected nodes. The training takes 100 epochs, which suffices for quick convergence.

(Table A2), and latitude, longitude, month, and rising/settling flag. The output parameters are specific humidity at six pressure levels between 975 and 850 hPa with a cadence of 25 hPa.

Compared to some earlier ML models (e.g., random forest, gradient boosting), CNN also learns the vertical cross-correlation within the 52-layer input SNR profiles, as well as within the targeted six layers of specific humidity profiles. We conducted a comprehensive search of best hyperparameters using the root-mean-square error (RMSE) as the loss function.

In the prediction step, 30 predictions were carried out given each input set of variables, the mean and standard deviation of which were used as the final prediction and estimated uncertainty. It is worth highlighting that in each convolutional and fully connected layer, a dropout rate of 0.25 is applied to generate the variation, which is then used to calculate

the standard deviation of the “ensemble prediction” as a way to measure the retrieval uncertainty. This so-called “Monte Carlo” dropout method was designed in ML as a standard technique to regularize model over-fitting (Srivastava et al., 2013) but was also employed widely as a Bayesian approximation to quantify model uncertainties (Gal and Ghahramani, 2016). Admittedly, the current method only provides a quantification for ML model errors. There is no consideration of SNR measurement errors, nor propagation of the error to the final retrievals at this moment, although this is certainly a procedure that can be in place in future works.

We also tried some earlier ML models, e.g., random forest (RF), gradient boosting (GB), and support vector machine (SVM), from the scikit-learn library and one deep learning model multilayer perceptron (MLP) from the pytorch library. The model performances are actually very close in terms of evaluating the RMSE except for the SVM, the latter of which performed discernibly worse than the rest ML models. It is not a surprise finding as this is a relatively simple and straightforward task that ML models should handle easily, but this is not the case for the multi-variable linear regression type of logistic models (hence, it explains the poor performance of SVM). As the main focus of this paper is science and new information content embedded in SNR signals, we will not deviate attention by spending more time discussing these model results. The semi-transparency of RF and GB models is appreciated by us though. We compared the feature importance rankings with the findings of Wu et al. (2022) and find high consistencies (e.g., high ranking of SNR at $H_{SL} = -100$ km in the tropics, and SNR at $H_{SL} = -80$ km ranks the top in the polar region).

2.3 Underlying physical mechanisms

It is necessary to provide a summary of the underlying physics to emphasize the solid physical ground for this product, so that readers will not misunderstand this as a pure statistics-based ad hoc finding. The underlying physical mechanisms to explain the observed high correlation between MABL water vapor and the GNSS SNR signal remain

an active research area. Wu et al. (2022) articulated that the diffractive effect on the RO signal under the condition of limb sounding through a sharp MABL can extend the signal below the sharp edge of the obstacle with a limited depth.

Both diffractive and refractive processes are required to happen along the radio wave propagation to produce the RO signal at deep H_{SL} . Another example (Sokolovskiy et al., 2024) found enhancement of SNR when super-refraction happens. In reality, complex MABL can produce a mixed effect in the soundings from a combination of conditions that include normal bending, grazing reflection, super-refraction, ducting, or diffraction (Sokolovskiy et al., 2014). As a result, sophisticated physical radiation transfer models (e.g., radiogram, canonical transform) can in principle be used but at the expense of high computational costs and are hence impractical operationally. Moreover, the retrieval itself is essentially still an under-constraint problem, which commonly occurs because satellite retrievals and assumptions (no matter physically making sense or not) need to be made to fully constrain the physical model. As the quasi-linear relationship is preserved at a profile-by-profile level with larger noise compared to the monthly gridded and smoothed data (Fig. 3), and the height dependency of the regression coefficient is highly non-linear (Fig. 2), a ML model is simply the best choice to extract the signal.

3 Results

3.1 Retrieval performance evaluation

As seen in the first comparison, Figs. 6 and 7 showcase the statistical closeness to the 1 : 1 line and the resemblance of geographical distributions for the 3 independent testing months: January–March 2018, for COSMIC-1. All six pressure levels are compiled together to make Fig. 6 but would otherwise look extremely similar if plotting level by level. The only deviation from the 1 : 1 line occurs at very small specific humidity values (ERA-5 specific humidity $< 1 \text{ g kg}^{-1}$); i.e., very dry conditions normally occur at high latitudes.

Such a discrepancy reveals itself more clearly when we map out the percentage difference (Fig. 7b). The largest percentage differences indeed are found at polar regions as well as near the coastal lines, with SNR-ML-retrieved humidity tending to be larger than ERA-5. Note that to satisfy ducting or other diffraction conditions in order to use SNR signal at deep H_{SL} , the surface is required to be flat. Therefore, the discrepancies around the coastal line are believed to be related to issues with SNR-ML retrievals when topography starts to play a role. However, as we will show later in Fig. 10, ERA-5 indeed shows consistent dry bias at high latitudes compared to independent radiosonde measurements. So, the SNR-ML retrieval might produce results closer to the truth, as will be seen later as well. Moreover, one can vi-

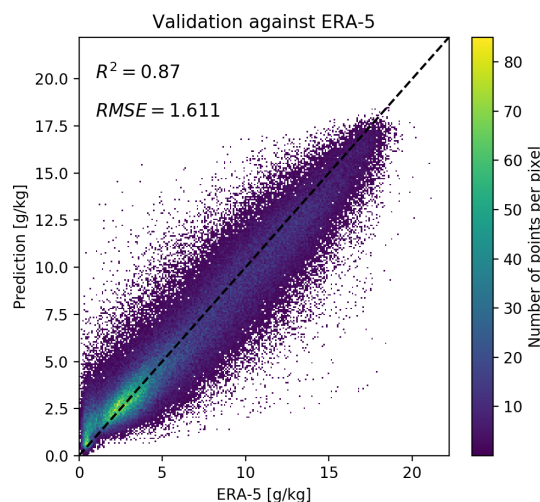


Figure 6. Heatmap for independent validation from January–March 2018 for COSMIC-1, combining all six levels together.

sually discern discrepancies in the tropical deep convection zone/Intertropical Convergence Zone (ITCZ), where ERA-5 in general is wetter than SNR-retrieved values. Such a discrepancy is not conspicuous in Fig. 7b, simply because of the large value in the denominator. We will also show later that none of the three datasets we will evaluate (SNR-ML retrieval, GNSS-RO wetPrf/wetPf2 retrieval, and ERA-5 re-analysis) capture the tropical MABL structures well. For the SNR-ML method, it is probably because the ducting assumption is easily and frequently violated in the tropical MABL.

3.2 Uncertainty quantification

Unfortunately, for very dry conditions, SNR-ML-method-retrieved specific humidity also inherently comes with large uncertainties, as can be clearly seen in Fig. 8. The SNR signal is too weak in this situation to yield any robust retrievals, even with powerful ML models. Although we still believe the SNR-ML retrievals might be “more correct” than ERA-5 for very dry conditions, in practice we mark any retrieval with greater than 50 % uncertainty with a quality flag in the published product, and those data do not pass quality control to be used later in this paper for independent validation or construction of the climatologies. This threshold only filters out about 2 % of the data with very weak SNR signals. If we were to apply a threshold of 20 %, about 16 % of data would be filtered out. In the later section when the diurnal cycle is compared using multi-year regional averaged data, we found that heavy averaging effectively beats down the noise, so as to reveal a visible diurnal signal in the extremely dry polar region, whereas ERA-5 is essentially a fixed value (Fig. 14b). We can also see from Fig. 8 that almost all SNR-ML retrievals greater than 2 g kg^{-1} pass quality control. Readers should keep in mind that our current uncertainty estimation

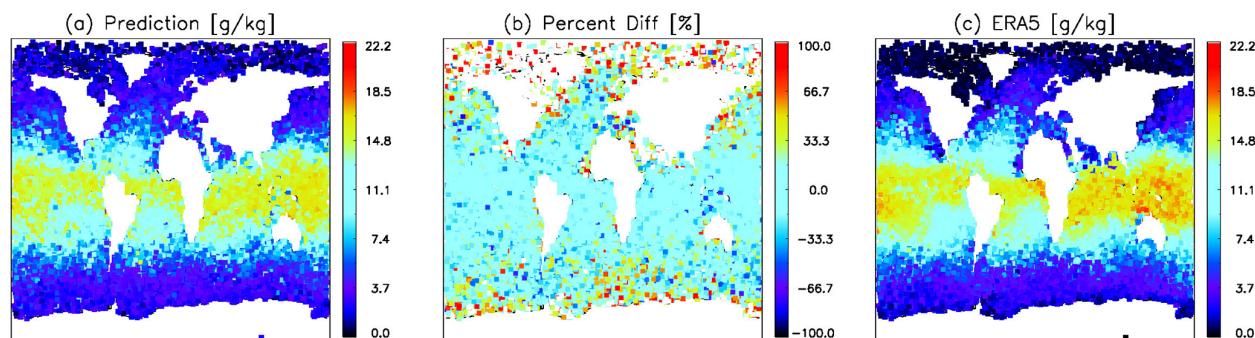


Figure 7. Geographic distribution of the (a) predicted values using COSMIC-1 SNR observations versus (c) ERA-5 validation values at 950 hPa for January–March, 2018. Panel (b) is the percentage difference between panels (a) and (c). Only ERA-5 samples that are collocated and coincident with COSMIC-1 SNR-ML retrievals are selected for this comparison.

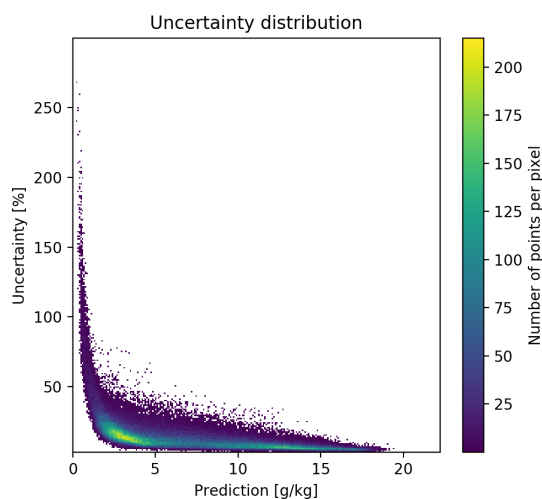


Figure 8. Percentage uncertainty distribution as a function of the predicted values.

approach underestimates the real uncertainty because it does not take SNR errors into account.

3.3 Comparison to independent radiosondes

In order to find collocation samples in every campaign, the collocation criteria are slightly different given the consideration of (1) the abundance of radiosonde/dropsonde profiles; (2) the typical spatial and temporal homogeneity of the local weather regime; and (3) the availability of daily COSMIC-1, COSMIC-2, Metop-A, and Metop-B profiles. In practice, for EUREC4A and ATOMIC, collocation is defined as longitude difference within 2° , latitude difference within 1.5° , and time difference within 1 h. For the Southern Ocean campaign, the thresholds become 4° , 2.5° , and 2 h correspondingly. For ARRecon and MAGIC campaigns, the thresholds are 4° , 1.5° , and 2 h.

Figure 9 shows the level-by-level comparison for all collocated samples from all campaigns. SNR-ML retrieval re-

sults are shown by filled color symbols, while wetPrf/wetPf2 retrievals are shown by open symbols. In addition, the averages from each campaign collocation subsets are connected together for better visual comparison against the 1 : 1 lines (solid black lines for SNR-ML retrieval and dotted black lines for wetPrf/wetPf2 retrievals). We can see that both SNR-ML retrievals and wetPrf/wetPf2 retrievals demonstrate generally good agreement with ground “truth” for different weather regimes. For the SNR-ML retrieval results, better correlations are found for the Southern Ocean (MARCUS campaign) and stratocumulus weather regimes (MAGIC campaign). Although wetPrf/wetPf2 results are highly comparable to the SNR-ML retrievals, the collocation samples are much sparser for the former (Table 3). This could be attributed to the frequent occurrence of super refractions in the stratocumulus region that causes a sampling bias of the wetPrf/wetPf2 results (Xie et al., 2010; Feng et al., 2020). Spreads are slightly larger during the atmospheric river events (ARRecons). SNR-ML retrievals show an overall better agreement compared to the wetPrf/wetPf2 retrievals at all six pressure levels, especially for the few extremely large specific humidity values ($> 12 \text{ g kg}^{-1}$). The means of all ARRecon collocated samples also suggest that SNR-ML retrieval is the only one that does not produce a bias, while wetPrf and wetPf2 are moderately (slightly) dry-biased in atmospheric river scenarios at > 900 (< 900) hPa. ERA-5 from each campaign (only considering samples that SNR-ML retrieval collocation is found) exhibits good agreement with the ground truth too. For the two deep tropics campaigns ATOMIC and EUREC4A, we can clearly see that none of the three datasets capture the humidity conditions in the MABL very well. They are all dry-biased, and ERA-5 reanalysis is slightly less dry-biased than GNSS retrieved values at 975 and 950 hPa. The SNR-ML method achieves overall comparable performance to ERA-5, which is expected because the model is trained on ERA-5.

For convenience in pinpointing ERA-5 MABL issues, we also make Fig. 10 as each valid radiosonde/dropsonde pro-

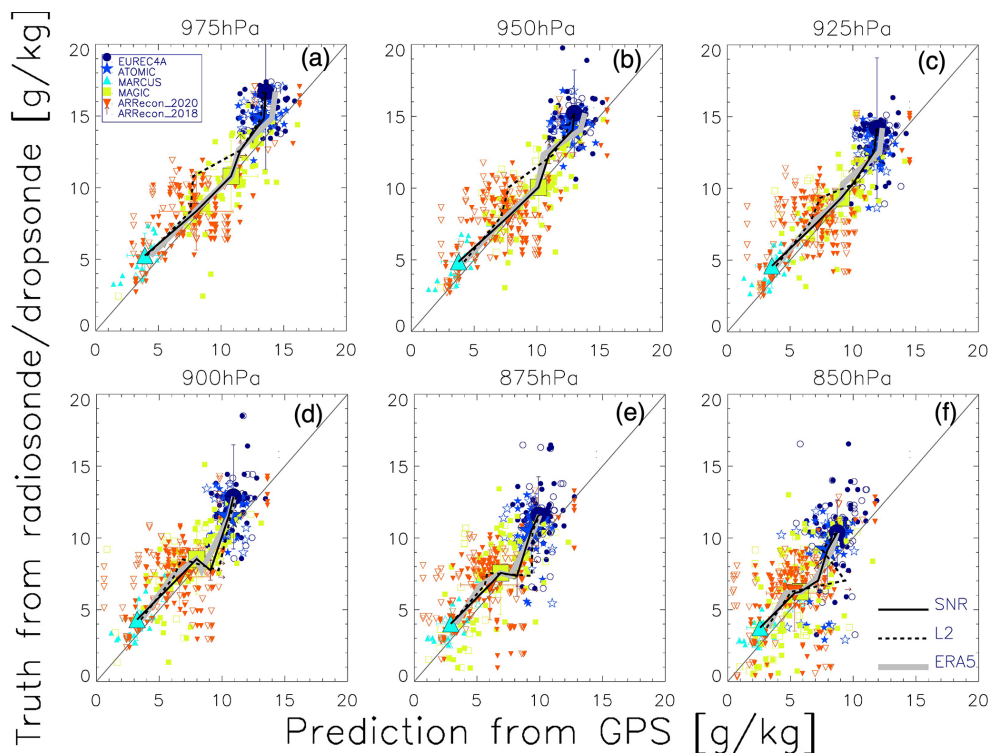


Figure 9. Scatter plots of collocated specific humidity [g kg^{-1}] comparison between radiosonde “truth” and retrievals from SNR (closed symbols) and wetPrf/wetPf2 standard retrieval (open symbols) for each pressure level. Thin black diagonal lines are the 1 : 1 lines for reference. The mean and standard deviation from the SNR-ML retrieval from each campaign are shown as bigger same-color symbols with black boundaries. In addition, these mean retrieved values from each campaign are connected by the bold black lines for SNR-ML retrievals, bold dash-dotted black lines for wetPrf/wetPf2 retrievals, and bold solid gray lines for ERA-5 from the subset where collocations are found for SNR-ML and radiosonde data samples.

Table 3. Number of collocated GNSS-radiosonde/dropsonde samples in each campaign. Two numbers in each cell are from SNR-ML method and wetPrf/wetPf2 product, respectively, and their percentage differences are shown in the parentheses.

Campaign name	975 hPa	950 hPa	925 hPa	900 hPa	875 hPa	850 hPa
EUREC4A	50, 19 (160 %)	50, 23 (117 %)	51, 29 (76 %)	51, 31 (65 %)	51, 34 (50 %)	51, 38 (34 %)
ATOMIC	49, 23 (113 %)	49, 27 (81 %)	49, 29 (69 %)	49, 29 (69 %)	49, 35 (40 %)	49, 44 (11 %)
MARCUS	13, 5 (160 %)	13, 7 (86 %)	13, 7 (86 %)	13, 7 (86 %)	13, 8 (63 %)	13, 9 (44 %)
MAGIC	72, 9 (700 %)	72, 25 (188 %)	72, 34 (112 %)	72, 40 (80 %)	72, 43 (67 %)	72, 46 (57 %)
ARRecon	120, 84 (43 %)	120, 101 (19 %)	120, 101 (19 %)	120, 106 (13 %)	120, 106 (13 %)	120, 106 (13 %)

file from all six campaigns can always be collocated with an ERA-5 reanalysis data sample within 1.5° longitude, 1° latitude, and 1 h difference. Now we can clearly see that ERA-5 frequently fails to produce the large variations in humidity in the trade-cumulus region (EUREC4A), the former of which tends to be always too wet. Otherwise, ERA-5 matches better than SNR-ML retrievals and wetPrf/wetPf2 retrievals in the deep tropics (EUREC4A and ATOMIC); however all of the three datasets contain persistent dry biases, as can also be seen in Fig. 9. Another discernible bias happens in the Southern Ocean during the MARCUS campaign, where ERA-5

is consistently dry-biased when specific humidity is below $\sim 3 \text{ g kg}^{-1}$. The subset used to make the gray lines in Fig. 9 is overlaid with open symbols, so we can make a straightforward and fair comparison between ERA-5 and SNR-ML retrievals. We can see that SNR-ML performs slightly better than ERA-5 in the atmospheric river scenarios (two AR-Recon campaigns) and slightly worse than ERA-5 in the stratocumulus region (MAGIC campaign), both of which are reflected in the correlation coefficient comparisons shown in Fig. 11 as well. Overall, ERA-5 shows a small dry bias globally at all levels, which agrees with early findings by

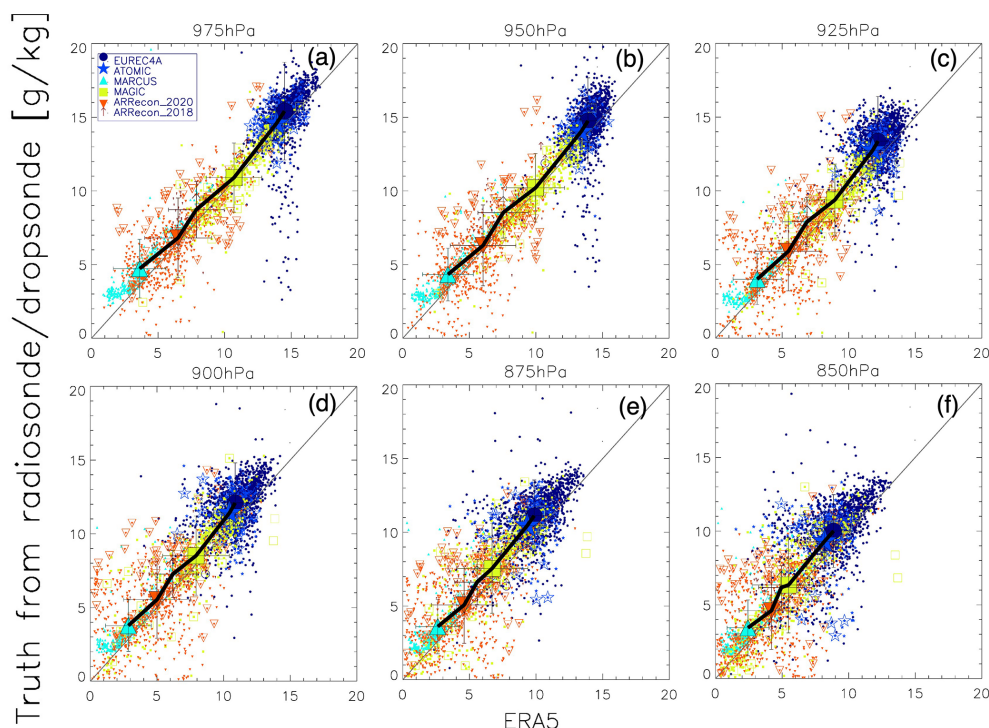


Figure 10. Same as Fig. 9 except for all available radiosonde/dropsonde samples in all these campaigns with collocated ERA-5 specific humidity. The means of each campaign are shown as bigger same-color symbols with black boundaries and standard deviation. The bold solid black line connects the mean values from each campaign.

Johnston et al. (2021), who used wetPf2 GNSS-RO retrievals to identify such a dry bias. Note that some of the campaign profiles (e.g., ARRecon dropsondes) are actually assimilated in the ERA-5 data, so it is not a completely independent validation strictly speaking. However, it is also worth noting that some previous publications employed ARRecon and EUREC4A radiosonde data as “ground truth” for evaluating ERA5 accuracy in capturing water vapor variabilities in the PBL (e.g., Cobb et al., 2021; Krüger et al., 2022).

The violin plots in Fig. 11 and numbers of collocated sample statistics in Table 3 help disentangle the merits and caveats of SNR-ML retrievals from multi-dimensional statistical metrics. Only correlation coefficients of all collocated samples collected from each campaign are displayed in Fig. 11. The ARRecon-2018 and ARRecon-2020 samples are further combined. From Fig. 11a, we can see again that the MABL specific humidity is not well captured in the tropics by either of the three datasets (EUREC4A and ATOMIC), but SNR-ML retrievals perform slightly better than the operational wetPf2 products in the deep tropics and trade-cumulus regions. In the rest of the three campaigns in the mid-latitudes and high latitudes, they all agree very well with the radiosonde/dropsonde ground truths. ERA-5 reanalysis does the best job in the high-latitude Southern Ocean (MARCUS) as well as the stratocumulus region (MAGIC), while in the atmospheric river regime, SNR-ML retrievals

outperform the wetPrf/wetPf2 retrievals as well as the ERA-5 reanalysis. It is worth noting that SNR-ML retrievals perform slightly better than wetPrf/wetPf2 retrievals in the stratocumulus region (MAGIC) in both the medians and the top-heavy skewness of its distributions, which can partially be attributed to the scarcity of wetPrf/wetPf2 collocation samples in this weather regime and known bias in the Level 2-retrieved refractivity gradient (Xie et al., 2010). For the polar region (MARCUS), although SNR-ML retrievals exhibit the lowest correlations among the three datasets albeit all correlations are statistically significant, it is inconclusive at this point to say that the SNR-ML method is not suitable for the polar region. As a matter of fact, the SNR-ML method generates the largest variabilities among the three when the PBL is extremely dry (Fig. 6), but the SNR in this situation is generally too weak to generate a robust retrieval (i.e., uncertainty too large compared to the retrieved value). The retrievals from the SNR-ML method in dry polar winters have more potential (see, e.g., Fig. 14) if future GNSS missions could improve the SNR.

Figure 11b demonstrates the robustness of the SNR-ML retrievals across all six PBL pressure levels. Although the highest positive correlations are always identified in ERA-5 and/or wetPrf/wetPf2 products, the medians of SNR-ML retrievals are consistently the highest with consistent top-heavy distribution except for 850 hPa, meaning that SNR-ML re-

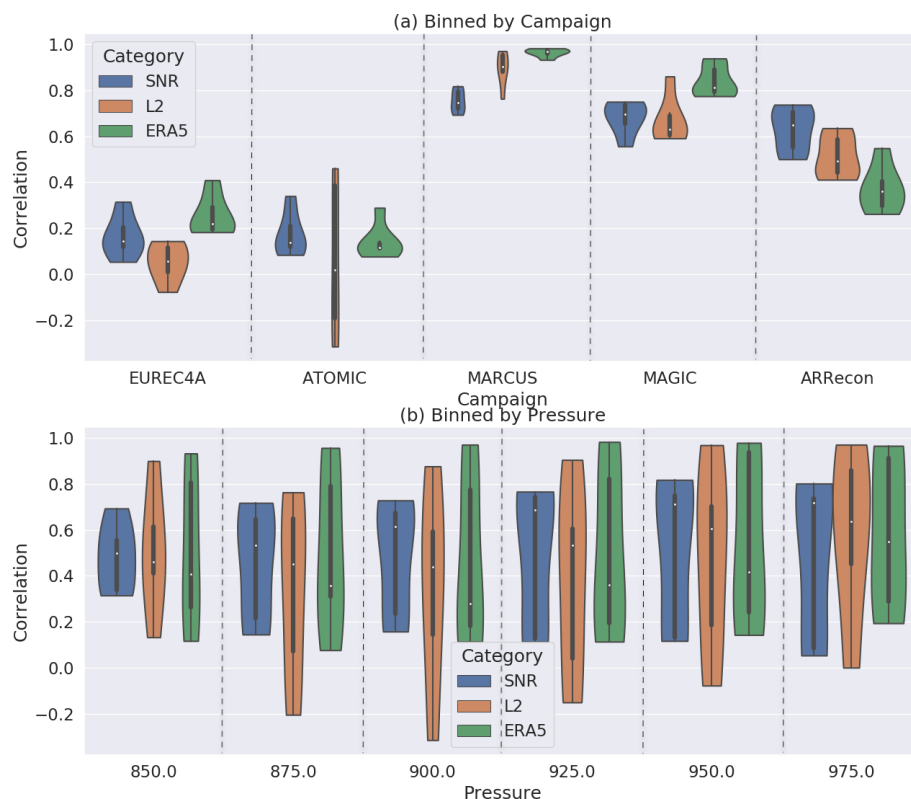


Figure 11. Violin plots of the correlation coefficients calculated from collocated samples for SNR-ML retrievals (blue), Level 2 retrievals (orange), and ERA-5 (green). Panel (a) is all-level statistics for each campaign, and panel (b) is all campaigns but binned by different pressure levels. Median, standard deviation, and minimum/maximum values and the skewness of the distribution are shown as the white dots, black box, extended vertical thin lines, and the horizontal widths in each violin, respectively. The number of total samples is listed above each violin. For ERA-5, only the subset of samples for which SNR-ML retrieval collocations are available are selected to calculate the statistics.

trievals agree with radiosonde/dropsonde “truths” more consistently, while ERA-5 and wetPrf/wetPf2 have more variation across different weather regimes. Of course all these conclusions are limited by the small collocation samples (309 in total), and we for sure need more extensive evaluation for this research product before mass production.

Another big advantage of the SNR-ML retrieval is its consistently higher success rate in the MABL compared with the wetPrf/wetPf2 product. This is clearly seen in Table 3, where the percentage difference between the two is listed in parentheses for each campaign at each pressure level. For the stratocumulus region (MAGIC campaign), when ducting or super-refraction happens frequently, the success rate of the SNR-ML method can go up to 700 % more than using the wetPrf product at the lowermost altitude. Although the superiority of the success rate of the SNR-ML retrievals gradually vanishes when getting closer to the MABL top, they are still more across the board than wetPrf/wetPf2 products.

To summarize the major findings for comparisons against the limited independent radiosonde/dropsonde datasets available over the open ocean, we can draw the following conclusions. Firstly, the quality of the SNR-ML retrievals is com-

parable to ERA-5 and the operational wetPrf/wetPf2 product. In atmospheric river weather regime, the SNR-ML method even outperforms the other two. The robustness and stable performance of SNR-ML retrievals remain the best within the MABL, although its advantage gradually vanishes with increasing height. Secondly, compared to the operational retrievals, the SNR-ML method can achieve 10 %–700 % more samples in the MABL, especially over stratocumulus regions, where ducting and super-refraction frequently occur, causing failure of operational retrievals. This suggests some unique value that the SNR-ML method can bring to the science community in facilitating understanding of the water vapor–stratocumulus coupling mechanisms. Although some of the “independent validation dataset” is not completely independent as the data may have been assimilated in the ERA-5, the fact that SNR-ML retrieval statistics outperform ERA-5 at all six pressure levels in diverse weather regimes proves that real physical information from SNR observations is learned and kept by the ML model for prediction; admittedly it is impossible to quantify how much the real observed information contributes without accurate physics-based model simulations.

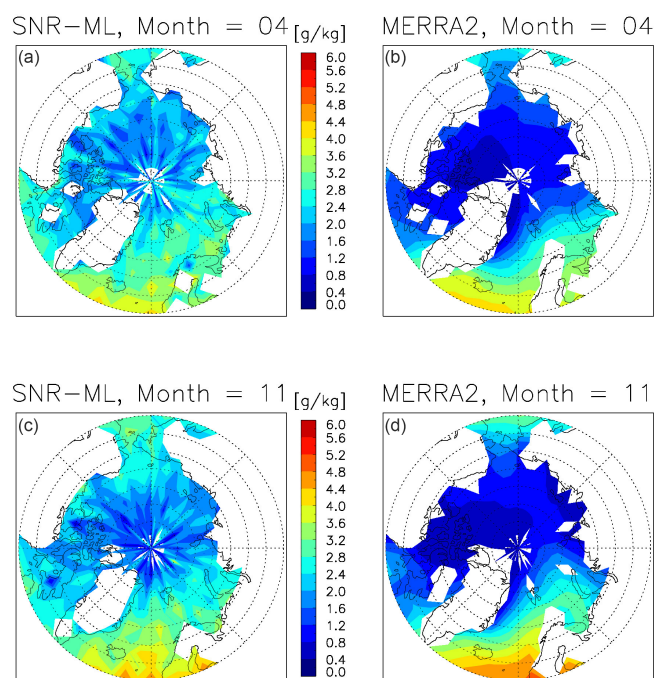


Figure 12. Monthly averages of 950 hPa specific humidity from COSMIC-1 SNR retrieval (a, c) compared to MERRA-2 reanalysis (b, d) for the Arctic during April (a, b) and November (c, d), 2012 and 2013.

4 Discussions

In this section, we present and discuss some use case examples in order to demonstrate how to use this SNR-ML MPBL specific humidity product to identify and even quantify model or reanalysis issues.

4.1 Climatology

Several previous studies suggest that MERRA-2 reanalysis has larger dry biases in the polar regions compared to ERA-5 (Johnston et al., 2021; Ganeshan and Yang, 2019), while some other studies using in situ campaign data suggested smaller dry bias in the MERRA-2 reanalysis (e.g., Seethala et al., 2021). Here we map out the climatological distribution of specific humidity retrieved using the SNR-ML method to track down geographical discrepancies in the Arctic (Fig. 12) and Antarctic (Fig. 13) with respect to MERRA-2. The coldest months were not selected because of the concern that the sea-ice-induced reflectometry signal might contaminate our SNR-ML retrieval results, but we did not exclude retrievals over possible glaciers for which MERRA-2 does not produce a valid value at 925 hPa because we used a fixed-terrain map. Therefore, direct comparison should not be considered wherever MERRA-2 value is blank.

Overall, again we can see the SNR-ML-method-retrieved polar MABL is much more humid than that from MERRA-2 in the Arctic during early spring and late fall seasons

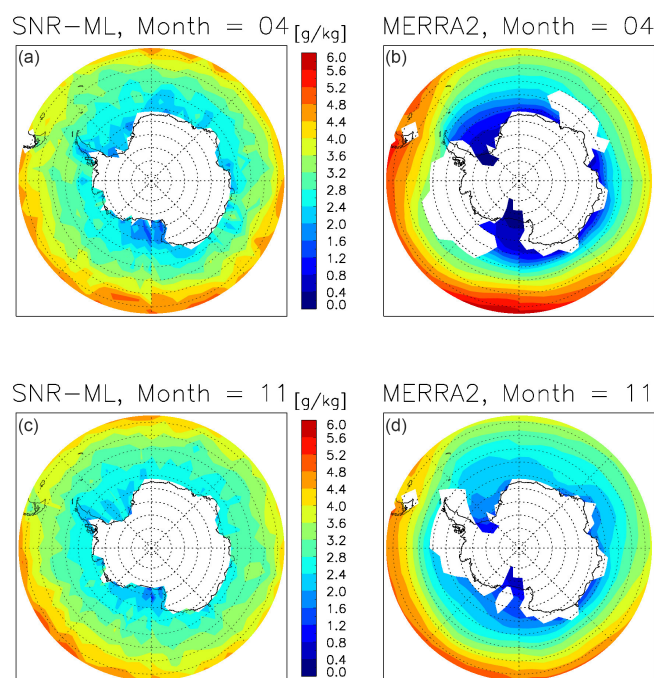


Figure 13. Same as Fig. 12 except for the Antarctic/Southern Ocean.

(> 100 % in most areas). If we neglect sampling-induced geographical inhomogeneities in the SNR-ML retrievals, we can actually see in Fig. 12 that the geographic distribution of highs and lows and their gradients are in general agreeable. The largest difference is that the wet intrusion along the Bering Strait seems to be too weak during both April and November in MERRA-2, which could account for the dry bias in the deep Arctic Ocean. Meanwhile, the wet intrusion associated with the North Atlantic overturning circulation seems to be too strong during November in MERRA-2. These discrepancies connect possible root causes down to the ocean circulation, and up to the Arctic front, and should be further investigated from a whole Earth system point of view.

Although the Southern Ocean and South Pole seem lacking in geographical variations (Fig. 13), we can actually observe some interesting potential issues related to topographies. For example, the tip of the Andes mountains effectively blocks MABL water transport across the mountains, but such a local effect on humidity appears further downstream in MERRA-2. The gradient of water vapor amount from north to south is apparently much weaker compared to MERRA-2, which impacts the latent heat and sensible heat flux quantification when considering global energy transport.

4.2 Diurnal variation

It is well known that global climate models (GCMs) have serious issues in reproducing cloud, precipitation, and convection diurnal cycles (e.g., Tian et al., 2004;

Yin and Porporato, 2017). Although such a problem is mostly attributed to the issues with cumulus parameterization schemes, we argue that the diurnal cycle of MABL water vapor also plays a nontrivial role as it ties closely to the shallow cumulus and stratocumulus; the latter, for example, is also closely related to the MABL height diurnal variation (e.g., Liu and Liang, 2010; Chepfer et al., 2019; Teixeira et al., 2021). Ground truth of the diurnal variation of MABL water vapor structures is extremely rare, probably because of the high cost associated with long-duration shipborne campaigns that often only launch radiosondes twice daily and hence cannot capture the diurnal variabilities. Therefore, here we only aim to show the discrepancies between ERA-5 and our SNR-ML-retrieval-generated diurnal cycle rather than determine which is right or which is wrong (Fig. 14).

In addition to the Southern Ocean MARCUS campaign and the atmospheric river regime ARRecon campaign that we have ground truths to compare with, several additional campaign regions and corresponding months are selected motivated by the observed diurnal variations of the MABL height established in Liu and Liang (2010). These two additional regions include the South Indian Ocean (INDOEX campaign, representing deep tropics) and the Arctic open ocean, representing polar winter conditions. The last one was added for the sole interest of checking if there is any diurnal cycle in the coldest season.

The averaged specific humidity at 875 hPa agrees well between the two datasets in the MARCUS and ARRecon campaigns, but the diurnal cycles in ERA-5 are too weak compared to the ground truths (red asterisks), while the diurnal cycles in the SNR-ML method are stronger. It is worth noting that neither SNR-ML nor ERA-5 reproduced a strong peak below 900 hPa around 10:00 am local time that both MARCUS and SOCRATES campaigns observed. The latter is another research campaign in the vicinity of MARCUS ship routes and season (Vömel and Brown, 2018) but was not employed for independent validation because of a lack of collocations with GNSS observations. This peak is probably associated with the shallow mixed-phase cloud pocket precipitation that is spatially so small and inhomogeneous in scale (D'Alessandro et al., 2021) that neither GNSS nor ERA-5 is able to capture or reproduce. The underestimation of the diurnal variability in the ARRecon campaign region is probably associated with the sampling bias because the campaign “truth” was sampled only during Atmospheric River (AR) events, while SNR-ML and ERA-5 sample the climatology background.

Although we have no ground truth to assess the diurnal cycles of MABL humidity in other two regions, we can tell that ERA-5 is wetter in the South Indian Ocean and significantly drier in the Arctic Ocean. Compared to the SNR-ML-method-retrieved diurnal cycle, the MABL water vapor diurnal cycle in ERA-5 is too weak in three areas but not the INDOEX campaign region. To put this in the context of the diurnal cycle of PBL height (Liu and Liang, 2010), in the

INDOEX campaign region, the diurnal cycle from ERA-5 and the SNR-ML method agrees reasonably well; both are anti-correlated with the diurnal cycle of PBL height change observed during that campaign. In the Arctic Ocean, ERA-5 has apparently set some arbitrary threshold to keep the water vapor at a constant low level, while SNR-ML retrievals suggest a weak diurnal variation.

Overall, we can see that the diurnal coupling between MABL water vapor, PBL height, and clouds is vastly different from area to area. However, ERA-5 likely under-produces the diurnal cycle amplitude of MABL water vapor globally. For SNR-ML retrievals, day-to-day variability often overwhelms the signal of diurnal cycle, yet the amplitude of diurnal cycle is still stronger and matches better with the limited ground truth. Ultimately, the lack of MABL water vapor ground “truth” measurements will continuously make observing and verifying the true diurnal cycle difficult. Other shipborne measurements, e.g., from upward pointing radiometers, might be helpful to disentangle this mystery in the future.

5 Conclusions

Marine planetary boundary layer (MABL) water vapor amount and vertical gradient are among the key factors to couple the ocean and atmosphere cloud, precipitation, and convection together, but meanwhile it is also among the hardest objects to retrieve from satellite remote sensing perspective. Given the penetration capability of GNSS signal through clouds, we proposed a novel way in Wu et al. (2022) to utilize the GNSS signal-to-noise ratio (SNR) in the deep H_{SL} to retrieve MABL water vapor profiles. In this paper, we demonstrated it is workable at a profile-by-profile level, leveraging the power of machine learning (ML) in capturing weak and non-linear signals. The surprising and novel findings in this paper are that the ML-trained model can make better predictions and outperform the training dataset (i.e., ERA-5) in some places, which demonstrates that real information content in the SNR signal is learned that would otherwise not be harnessed using traditional statistical methods. The new SNR-ML retrieval has a more stable performance compared to the operational wetPrf/wetPf2 GNSS-RO retrievals, and it can produce 20 %–700 % more successful retrievals in the lowest 1 km, where observations are critical to understand ocean–atmosphere exchange.

We then showed two use cases to demonstrate possible ways to use this dataset. There are no conclusive results because of lack of ground “truth” to validate, but we do find both reanalyses tend to systematically produce dry biases at high latitudes and diurnal cycles that are too weak over global oceans. This SNR-ML retrieval dataset also has its own caveats. Whenever the “ducting” condition is violated (e.g., coastal topography, convective tower, mixing, and turbulence in the MABL), the fundamental assumption breaks

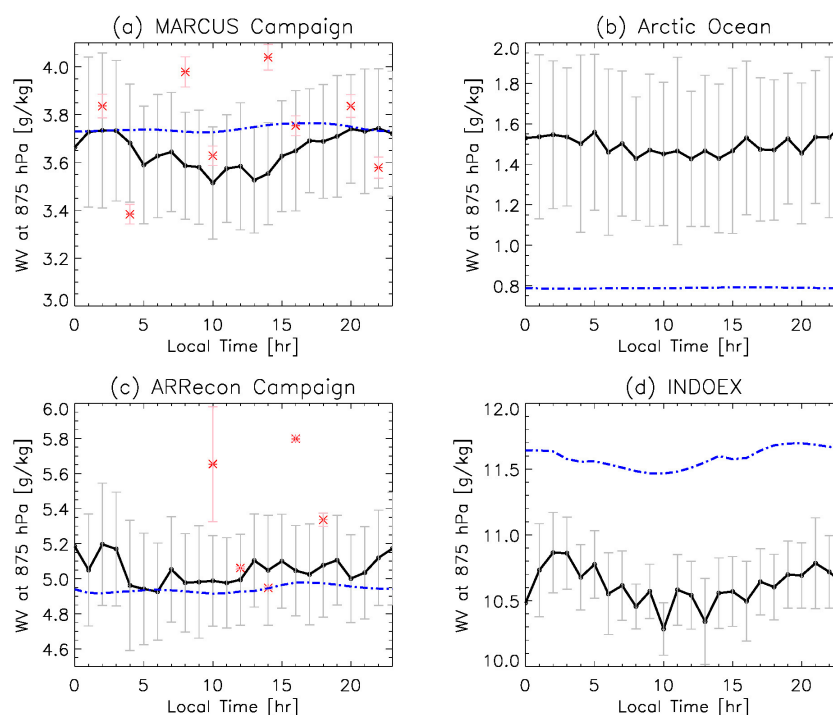


Figure 14. Multi-year mean diurnal variation of 875 hPa specific humidity retrieved from all four missions (black with error bars in gray) and from ERA-5 hourly reanalysis (dash-dotted blue) during November–March for **(a)** the MARCUS campaign region, 60–150° E, 60–40° S; **(b)** the Arctic Ocean, 180° W–180° E, 70–90° N; **(c)** the ARRecon campaign region, 160–120° W, 20–50° N; and **(d)** the INDOEX campaign region, 55–75° E, 25–15° S. The MARCUS radiosonde and ARRecon dropsonde “truths” are overlaid in panels **(a)** and **(c)** as asterisks with standard deviations shown in pink vertical bars.

down, resulting in poor performance. More extensive comparisons and validations against other high-quality ground measurements are needed in the future.

Based on results from this work, one can see that deep SNR can complement the current GNSS-RO operational bending angle product for retrieving PBL information for different PBL conditions. A merged product is certainly of interest to future investigations, but fully understanding the physical mechanisms behind the reemerged deep SNR signal is the foundation for other downstream applications (e.g., data assimilation). Right now this can be considered a stand-alone observational product for independent comparison or validation against model simulations or other observations.

Appendix A

Table A1. Summary of GNSS-RO instrument noise (σ) used in this work, separated by rising and setting modes.

Instrument name	Orbit	Noise (σ)
COSMIC-1/C1	Rising	10.1
	Setting	10.9
COSMIC-1/C2	Rising	10.2
	Setting	10.9
COSMIC-1/C3	Rising	9.6
	Setting	10.4
COSMIC-1/C4	Rising	10.6
	Setting	11.2
COSMIC-1/C5	Rising	10.1
	Setting	11.1
COSMIC-1/C6	Rising	9.2
	Setting	10.7
COSMIC-2/E1	Rising	17.0
	Setting	17.5
COSMIC-2/E2	Rising	17.5
	Setting	17.8
COSMIC-2/E3	Rising	17.2
	Setting	17.9
COSMIC-2/E4	Rising	17.5
	Setting	17.7
COSMIC-2/E5	Rising	17.4
	Setting	17.8
COSMIC-2/E6	Rising	17.5
	Setting	17.8

Table A2. Excess phase L1 grid for this work.

Parameter	Grid values
$\text{Log}_{10}(\phi_{L1})$	1.26245, 1.33846, 1.41162, 1.48144, 1.54777, 1.62428, 1.69679, 1.76530, 1.82995, 1.89098, 1.94866, 1.97000, 2.00325, 2.02000, 2.05500, 2.08000, 2.10415, 2.13000, 2.15091, 2.17000, 2.19548, 2.23805, 2.27875, 2.30103, 2.32222, 2.37000, 2.41497, 2.44000, 2.55630, 2.59000, 2.63000, 2.69020, 2.75000, 2.81291, 2.86000, 2.92428, 2.95000, 3.02531, 3.10000, 3.11727, 3.15000, 3.20140, 3.22000, 3.25000, 3.27875, 3.30000, 3.32000, 3.35025, 3.41664, 3.47857, 3.53656, 3.59106
Rough corresponding H_{SL} [km]	19, 17, 15, 13, 11, 10, 9, 8, 7, 6, 5, 4, 3, 2, 1, -1, -2, -3, -5, -7, -9, -11, -13, -15, -17, -19, -20, -23, -26, -30, -33, -37, -40, -50, -60, -70, -80, -90, -92, -94, -96, -98, -100, -102, -104, -106, -108, -110, -120, -130, -140, -150

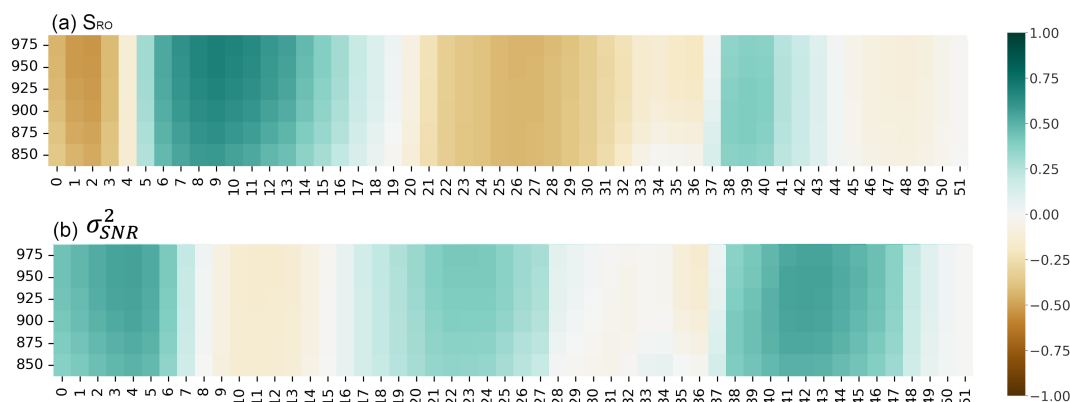


Figure A1. Same as Fig. 2 except for the Metop-A training dataset.

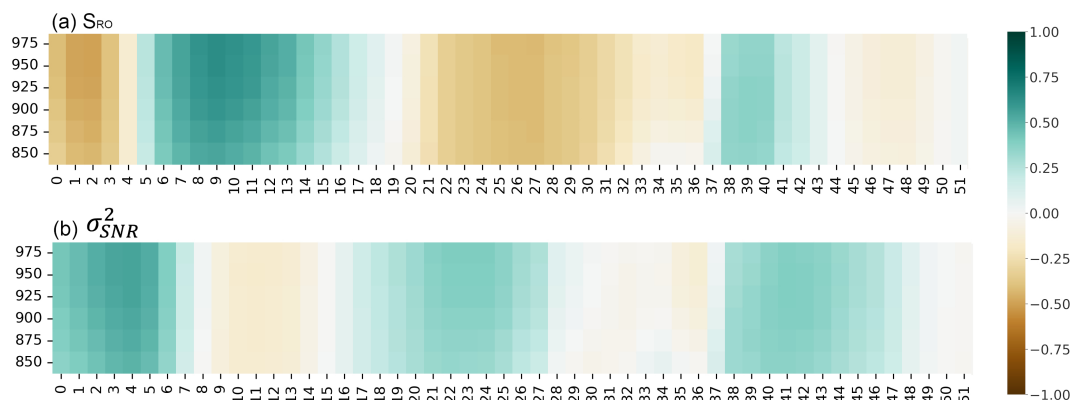


Figure A2. Same as Fig. 2 except for the Metop-B training dataset.

Data availability. The Level 2 SNR-ML retrieval product for the prediction period (see Table 1) has been published on Zenodo (<https://doi.org/10.5281/zenodo.13946112>; Gong and Dong, 2024). We welcome use and feedback.

COSMIC-1 and COSMIC-2 Level-1 and Level 2 data are downloaded from <https://doi.org/10.5065/8r12-hs65> (Sokolovskiy, 2020). Metop-A and Metop-B data are downloaded from <https://gpsmet.umd.edu/gnssro/download.php> (last access: 21 August 2025). ATOMIC data are downloaded from <https://psl.noaa.gov/atomic/data/> (last access: 21 August 2025). EU-REC4A data are downloaded from <https://doi.org/10.25326/137> (Stephan et al., 2025). SOCRATES data are downloaded from <https://doi.org/10.5065/D69P30HG> (NCAR, 2025). MARCUS data and MAGIC data are downloaded from the ARM data request portal (<https://adc.arm.gov/discovery/#/>, last access: 21 August 2025). ARRecon data are downloaded from https://cw3e.ucsd.edu/arrecon_data/ (last access: 21 August 2025) and specially processed to fit the needs of this research. Interested users are encouraged to contact the last author for assistance of post-processed data.

Author contributions. DLW came up with the initial idea. JG designed the methodology, executed the plan, built the model, and conducted the validation and data analysis. MB helped conduct the hyperparameter grid search. MG provided Fig. 1. MZ provided the high-vertical-resolution ARRecon data. All authors participated in result discussion and interpretation.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: 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.

Special issue statement. This article is part of the special issue "Observing atmosphere and climate with occultation techniques –

results from the OPAC-IROWG 2022 workshop”. It is a result of the International Workshop on Occultations for Probing Atmosphere and Climate, Leibnitz, Austria, 8–14 September 2022.

Acknowledgements. Jie Gong is grateful to Mariel Friburg at NASA Goddard for providing financial support for Michelle Badalov. We thank the editor and two reviewers for their thorough and insightful comments, which helped greatly in improving the readability and clarity of this paper.

Financial support. This research has been supported by the National Aeronautics and Space Administration (grant no. 19-GNSS19-0005). Publication cost is covered by NASA (grant no. DSI-QRS-24-0001). Minghua Zheng received funding from ONR (grant no. N00014-24-1-2698).

Review statement. This paper was edited by C. Marquardt and reviewed by two anonymous referees.

References

- Boisvert, L. N., Wu, D. L., Vihma, T., Susskind, J.: Verification of air/surface humidity differences from AIRS and ERA-Interim in support of turbulent flux estimation in the Arctic, *J. Geophys. Res.*, 120, 945–963, <https://doi.org/10.1002/2014JD021666>, 2015.
- Chang, H., Lee, J., Yoon, H., Morton, Y. J., and Saltman, A.: Performance assessment of radio occultation data from GeoOptics by comparing with COSMIC data, *Earth Planets Space*, 74, 108, <https://doi.org/10.1186/s40623-022-01667-6>, 2022.
- Chepfer, H., Brogniez, H., and Noel, V.: Diurnal variations of cloud and relative humidity profiles across the tropics, *Sci. Rep.*, 9, 16045, <https://doi.org/10.1038/s41598-019-52437-6>, 2019.
- Cobb, A., Michaelis, A., Iacobellis, S., Ralph, F. M., and Delle Monache, L.: Atmospheric River Sectors: definition and characteristics observed using dropsondes from 2014–20 CalWater and AR Recon, *Mon. Weather Rev.*, 149, 623–644, <https://doi.org/10.1175/MWR-D-20-0177.1>, 2021.
- D’Alessandro, J. J., McFarquhar, G. M., Wu, W., Stith, J. L., Jensen, J. B., and Rauber, R. M.: Characterizing the Occurrence and Spatial Heterogeneity of Liquid, Ice, and Mixed Phase Low-Level Clouds Over the Southern Ocean Using in Situ Observations Acquired During SOCRATES, *J. Geophys. Res.*, 126, e2020JD034482, <https://doi.org/10.1029/2020JD034482>, 2021.
- Feng, X., Xie, F., Ao, C., and Anthes, R.: Ducting and Biases of GPS Radio Occultation Bending Angle and Refractivity in the Moist Lower Troposphere, *J. Atmos. Ocean. Tech.*, 37, 1013–1025, <https://doi.org/10.1175/JTECH-D-19-0206.1>, 2020.
- Gal, Y. and Ghahramani, Z.: Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning, *Proceedings of the 33rd International Conference on Machine Learning*, 48, <https://proceedings.mlr.press/v48/gal16.pdf> (last access: 21 August 2025), 2016.
- Ganeshan, M. and Yang, Y.: Evaluation of the Antarctic Boundary Layer Thermodynamic Structure in MERRA2 Using Dropsonde Observations from the Concordiasi Campaign, *Earth Space Sci.*, 6, 2397–2409, <https://doi.org/10.1029/2019EA000890>, 2019.
- Ganeshan, M., Wu, D. L., Santanello, J. A., Gong, J., Ao, C., Vergados, P., and Nelson, K. J.: Exploring commercial Global Navigation Satellite System (GNSS) radio occultation (RO) products for planetary boundary layer studies in the Arctic, *Atmos. Meas. Tech.*, 18, 1389–1403, <https://doi.org/10.5194/amt-18-1389-2025>, 2025.
- George, G., Stevens, B., Bony, S., Pincus, R., Fairall, C., Schulz, H., Kölling, T., Kalen, Q. T., Klingebiel, M., Konow, H., Lundry, A., Prange, M., and Radtke, J.: JOANNE: Joint dropsonde Observations of the Atmosphere in tropical North Atlantic meso-scale Environments, *Earth Syst. Sci. Data*, 13, 5253–5272, <https://doi.org/10.5194/essd-13-5253-2021>, 2021.
- Gong, J. and Dong, W.: GNSS deep SNR retrievals of marine atmosphere boundary layer (MABL) specific humidity, Zenodo [data set], <https://doi.org/10.5281/zenodo.13946112>, 2024.
- Johnston, B. R., Randel, W. J., and Sjöberg, J. P.: Evaluation of Tropospheric Moisture Characteristics Among COSMIC-2, ERA5 and MERRA-2 in the Tropics and Subtropics, *Remote Sens.*, 13, 880, <https://doi.org/10.3390/rs13050880>, 2021.
- Kapoor, S. and Narayanan, A.: Leakage and the reproducibility crisis in machine-learning-based science, *Patterns*, 4, 9, <https://doi.org/10.1016/j.patter.2023.100804>, 2023.
- Keeler, E., Burk, K., Kyrouac, J.: ARM – Balloon-borne sounding system (BBSS) WNP output data, sondewnpn.b1, ARM [data set], <https://doi.org/10.5439/1595321>, 2022.
- Krüger, K., Schäfler, A., Wirth, M., Weissmann, M., and Craig, G. C.: Vertical structure of the lower-stratospheric moist bias in the ERA5 reanalysis and its connection to mixing processes, *Atmos. Chem. Phys.*, 22, 15559–15577, <https://doi.org/10.5194/acp-22-15559-2022>, 2022.
- Kuo, Y., Sokolovskiy, S., Anthes, R. A., and Vandenberghe, F.: Assimilation of GPS Radio Occultation Data for Numerical Weather Prediction, *Space Sci.*, 11, 157–186, [https://doi.org/10.3319/TAO.2000.11.1.157\(COSMIC\)](https://doi.org/10.3319/TAO.2000.11.1.157(COSMIC)), 2000.
- LeCun, Y., Bengio, Y., and Hinton, G.: Deep learning, *Nature*, 521, 436–444, <https://doi.org/10.1038/nature14539>, 2015.
- Liu, S. and Liang, X.: Observed Diurnal Cycle Climatology of Planetary Boundary Layer Height, *J. Climate*, 23, 21, 5790–5809, <https://doi.org/10.1175/2010JCLI3552.1>, 2010.
- Millán, L. F., Lebsock, M. D., and Teixeira, J.: Variability of bulk water vapor content in the marine cloudy boundary layers from microwave and near-infrared imagery, *Atmos. Chem. Phys.*, 19, 8491–8502, <https://doi.org/10.5194/acp-19-8491-2019>, 2019.
- Milestein, A.: Planetary Boundary layer (PBL) Final Report, Lincoln Lab, MIT, PFR-10373, <https://apps.dtic.mil/sti/pdfs/AD1166514.pdf> (last access: 22 August 2025), 2022.
- Milestein, A. and Blackwell, B.: Neural network temperature and moisture retrieval algorithm validation for AIRS/AMSU and CrIS/ATMS, *J. Geophys. Res.-Atmos.*, 121, 1414–1430, <https://doi.org/10.1002/2015JD024008>, 2016.
- Milstein, A., Santanello, J. A., and Blackwell, B.: Detail Enhancement of AIRS/AMSU Temperature and Moisture Profiles Using a 3D Deep Neural Network, *Artificial Intelligence in Earth Science*, 2, e220037, <https://doi.org/10.1175/AIES-D-22-0037.1>, 2023.
- NCAR: NCAR/EOL ISS Radiosonde Data, NCAR [data set], <https://doi.org/10.5065/D69P30HG>, 2025.

- Seethala, C., Zuidema, P., Edson, J., Brunke, M., Chen, G., Li, X.-Y., Painemal, D., Robinson, C., Shingler, T., Shook, M., Sorooshian, A., Thornhill, L., Tornow, F., Wang, H., Zeng, X., and Ziemba, L.: On assessing ERA5 and MERRA2 representations of cold-air outbreaks across the Gulf Stream, *Geophys. Res. Lett.*, 48, e2021GL094364, <https://doi.org/10.1029/2021GL094364>, 2021.
- Sokolovskiy, S.: Standard RO Inversions in the Neutral Atmosphere 2013–2020 (Processing Steps and Explanation of Data), UCAR [data set], <https://doi.org/10.5065/8r12-hs65>, 2020.
- Sokolovskiy, S., Schreiner, W., Zeng, Z., Hunt, D., Lin, Y.-C., and Kuo, Y.-H.: Observation, analysis, and modeling of deep radio occultation signals: Effects of tropospheric ducts and interfering signals, *Radio Sci.*, 49, 954–970, <https://doi.org/10.1002/2014RS005436>, 2014.
- Sokolovskiy, S., Zheng, Z., Hunt, D. C., Weiss, J.-P., Braun, J. J., Schreiner, W. S., Anthes, R. A., Kuo, Y.-H., Zhang, H., Lenchow, D., and Vanhove, T.: Detection of Superrefraction at the Top of the Atmospheric Boundary Layer from COSMIC-2 Radio Occultations, *J. Atmos. Ocean. Tech.*, 41, 65–78, <https://doi.org/10.1175/JTECH-D-22-0100.1>, 2024.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting, *J. Mach. Learn. Res.*, 15, 1929–1958, <http://jmlr.org/papers/v15/srivastava14a.html> (last access: 22 August 2025), 2014.
- Stephan, C. C., Schnitt, S., Schulz, H., Bellenger, H., de Szoeko, S. P., Acquistapace, C., Baier, K., Dauhut, T., Laxenaire, R., Morfa-Avalos, Y., Person, R., Quiñones Meléndez, E., Bagheri, G., Böck, T., Daley, A., Güttler, J., Helfer, K. C., Los, S. A., Neuberger, A., Röttenbacher, J., Raeke, A., Ringel, M., Ritschel, M., Sadoulet, P., Schirmacher, I., Stolla, M. K., Wright, E., Charpentier, B., Doerenbecher, A., Wilson, R., Jansen, F., Kinne, S., Reverdin, G., Speich, S., Bony, S., and Stevens, B.: Ship- and island-based atmospheric soundings from the 2020 EUREC4A field campaign, *Earth Syst. Sci. Data*, 13, 491–514, <https://doi.org/10.5194/essd-13-491-2021>, 2021.
- Stephan, C. C., Schnitt, S., Schulz, H., and Belleger, H.: Radiosonde measurements from the EUREC4A field campaign (v3.0.0), EUREC4A [data set], <https://doi.org/10.25326/137>, 2025.
- Teixeira, J., Piepmeier, J. R., Nehrir, A. R., Ao, C. O., Chen, S. S., Clayson, C. A., Fridlind, A. M., Lebsock, M., McCarty, W., Salmun, M., Santanello, J. A., Turner, D. A., Wang, Z., and Zeng, X.: Toward A Global Planetary Boundary Layer Observing System: The NASA PBL Incubation Study Team Report, NASA, <https://ntrs.nasa.gov/api/citations/20230001633/downloads/AFridlindPBLTowardsReport.pdf> (last access: 22 August 2025), 2021.
- Tian, B., Soden, B., and Wu, X.: Diurnal cycle of convection, clouds, and water vapor in the tropical upper troposphere: Satellites versus a general circulation model, *J. Geophys. Res.*, 109, D10101, <https://doi.org/10.1029/2003JD004117>, 2004.
- Virman, M., Bister, M., Räisänen, J., Sinclair, V. A., and Järvinen, H.: Radiosonde comparison of ERA5 and ERA-Interim re-analysis datasets over tropical oceans, *Tellus A*, 73, 1929752, <https://doi.org/10.1080/16000870.2021.1929752>, 2021.
- Vömel, H. and Brown, W.: SOCRATES-2018 Radiosonde Data Quality Report, UCAR/NCAR – Earth Observing Laboratory, University Corporation for Atmospheric Research, <https://doi.org/10.5065/D69P30HG>, 2018.
- Wee, T.-K., Anthes, R. A., Hunt, D. C., Schreiner, W. S., and Kuo, Y.-H.: Atmospheric GNSS RO 1D-Var in Use at UCAR: Description and Validation, *Remote Sens.*, 14, 5614, <https://doi.org/10.3390/rs14215614>, 2022.
- Wu, D. L., Gong, J., and Ganeshan, M.: GNSS-RO Deep Refraction Signals from Moist Marine Atmospheric Boundary Layer (MABL), *Atmosphere*, 13, 6, 953, <https://doi.org/10.3390/atmos13060953>, 2022.
- Xie, F., Wu, D. L., Ao, C. O., Kursinski, E. R., Mannucci, A. J., and Syndergaard, S.: Super-refraction effects on GPS radio occultation refractivity in marine boundary layers, *Geophys. Res. Lett.*, 37, L11805, <https://doi.org/10.1029/2010GL043299>, 2010.
- Xu, X. and Zou, X.: COSMIC-2 RO Profile Ending at PBL Top with Strong Vertical Gradient of Refractivity, *J. Remote Sens.*, 14, 2189, <https://doi.org/10.3390/rs14092189>, 2022.
- Ye, J., Liu, L., Wang, Q., Hu, S., and Li, S.: A Novel Machine Learning Algorithm for Planetary Boundary Layer Height Estimation Using AERI Measurement Data, *IEEE Geosci.*, 19, 1–5, <https://doi.org/10.1109/LGRS.2021.3073048>, 2021.
- Yin, J. and Porporato, A.: Diurnal cloud cycle biases in climate models, *Nat. Commun.*, 8, 2269, <https://doi.org/10.1038/s41467-017-02369-4>, 2017.
- Zheng, M., Torn, R., Delle Monache, L., Doyle, J., Ralph, F. M., Tallapragada, V., Davis, C., Steinhoff, D., Wu, X., Wilson, A. M., Papadopoulos, C., and Mulrooney, P.: An Assessment of Dropsonde Sampling Strategies for Atmospheric River Reconnaissance, *Mon. Weather Rev.*, 152, 811–835, <https://doi.org/10.1175/MWR-D-23-0111.1>, 2024.
- Zhran, M.: An evaluation of GNSS radio occultation atmospheric profiles from Sentinel-6, *Egyptian J. Remote Sensing and Space Sci.*, 26, 654–665, <https://doi.org/10.1016/j.ejrs.2023.07.004>, 2023.