MiRS-AI: A Data-Driven Framework for Accurate and Efficient Atmospheric Profile Retrieval from Microwave Observations

Xingming Liang [1] **Yong-Keun Lee** [1] **Shuyan Liu** [1] **Xun (John) Yang** [1] **Wei Wang** [1] **Huan Meng** [2]

Abstract

Radiative transfer–based retrieval systems such as NOAA's Microwave Integrated Retrieval System (MiRS) are widely used to estimate atmospheric states from satellite observations but are computationally expensive and limited by the ill-posed nature of the inversion problem. We present MiRS-AI, a deep learning framework that formulates atmospheric retrieval as a data-driven surrogate modeling task. The framework includes a pixel-wise ResNet model that emulates traditional 1D-Var retrieval and a spatial UNet-Transformer (UNT) model that captures multi-scale dependencies across neighboring observations. Trained on global Advanced Technology Microwave Sounder (ATMS) observations and ECMWF analysis data, MiRS-AI achieves improved accuracy with lower bias and variance than operational MiRS while providing orders-of-magnitude speedup. The UNT model further improves performance in dynamically complex regions such as tropical cyclones, highlighting the importance of spatial context, though it may partially inherit smoothing from reference data and underrepresent fine-scale structures. These results demonstrate that data-driven surrogate models can effectively approximate physics-based retrieval systems, offering a scalable path for future extensions to additional atmospheric parameters and multimodal observations.

1. Introduction

Satellite-based atmospheric profile retrieval is a fundamental inverse problem in Earth system science, enabling applications in numerical weather prediction, climate monitoring, and extreme weather analysis. Operational systems such as the NOAA Microwave Integrated Retrieval System (MiRS, Boukabara et al., 2011; Lee et al., 2021) estimate atmospheric temperature (T) and water vapor (WV) profiles from passive microwave observations using a 1D variational (1D-Var) framework. In this formulation, retrieval is performed by minimizing a cost function that balances observational consistency with prior information through a physics-based radiative transfer model.

While MiRS provides physically consistent solutions, it faces two fundamental challenges. First, the inversion problem is inherently ill-posed, and retrieval accuracy depends strongly on prior assumptions such as background and observation error statistics, which may limit the ability to fully exploit observational information. Second, the computational cost of iterative optimization—driven by repeated evaluations of the radiative transfer model and its Jacobians—limits scalability for real-time and large-scale applications. These challenges motivate the development of efficient surrogate approaches that can approximate physics-based retrieval while improving computational efficiency.

In this work, we propose MiRS-AI, a deep learning–based framework that reformulates atmospheric retrieval as a data-driven surrogate modeling problem, directly learning the mapping from satellite observations to atmospheric state variables. The framework includes two complementary architectures:

- A pixel-wise ResNet model (Liang et al., 2026), which emulates the structure of traditional 1D-Var retrieval.
- A spatially aware UNet-Transformer (UNT) model, which captures multiscale spatial dependencies across neighboring observations.

By incorporating spatial context, MiRS-AI moves beyond the independent-pixel assumption inherent in 1D-Var systems and better represents complex atmospheric structures.

We train and evaluate MiRS-AI using globally distributed datasets constructed from the Advanced Technology Microwave Sounder (ATMS, Liang et al., 2022) observations and the European Centre for Medium-Range Weather Forecasts (ECMWF, Liang et al., 2022) analysis profiles. Results show that MiRS-AI achieves improved accuracy compared

[1]CISESS, University of Maryland, College Park, MD, USA. [2]STAR, NOAA/NESDIS, College Park, MD, USA. Correspondence to: Xingming Liang <xingming.laing@noaa.gov>.

Accepted at ICML 2026 AI for Science Workshop.

to MiRS while providing substantial computational speedup. In particular, the UNT model significantly enhances retrieval performance in dynamically complex regimes such as tropical cyclones, highlighting the importance of spatial modeling in geophysical inverse problems.

From a broader perspective, this work demonstrates that deep learning–based surrogate models can effectively approximate physics-based retrieval systems while capturing multiscale structure in atmospheric fields. The proposed framework provides a foundation for integration into higher-dimensional data assimilation systems (e.g., 3D/4D-Var, Wang et al., 2013) and can be extended to incorporate additional atmospheric parameters, such as cloud profiles. It can also be adapted to leverage additional microwave and hyperspectral sensors through dimensionality reduction techniques such as principal component analysis (PCA) combined with deep learning.

Contributions. This paper makes the following contributions:

- We introduce MiRS-AI, a unified deep learning framework for atmospheric profile retrieval that replaces iterative variational optimization with direct inference
- We propose a spatial retrieval architecture (UNT) that leverages multiscale contextual information to improve performance in complex atmospheric conditions.
- We identify a key limitation of UNT: its reliance on smoothed ECMWF analyses may underrepresent fine-scale thermodynamic structures, highlighting the need for observation-driven constraints or higher-resolution references to capture small-scale variability in tropical cyclones.
- We demonstrate, through global and extreme-weather evaluations, that MiRS-AI achieves higher accuracy and significant computational efficiency gains compared to the operational MiRS system.
- We highlight the potential of AI-based surrogate modeling as a scalable pathway toward real-time Earth system analysis.

2. Related Work and Method

Atmospheric profile retrieval from satellite observations is commonly formulated as a nonlinear inverse problem. In the operational MiRS system, this problem is solved using a variational framework, where the atmospheric state is estimated by minimizing a cost function that balances observational consistency and prior information:

$$\begin{aligned} J(x) =& (x - x_b)^T B^{-1} (x - x_b) \\ &+ (y - H(x))^T R^{-1} (y - H(x)) \end{aligned} \tag{1}$$

where x denotes the atmospheric state vector (e.g., temperature and water vapor profiles), x_b is the background state, B is the background error covariance matrix, y represents observed brightness temperatures, H(x) is the radiative transfer forward model, and R is the observation error covariance. Solving this optimization problem requires iterative updates, involving repeated evaluations of the forward model and its Jacobian, which constitutes the dominant computational cost in MiRS.

In this work, we reformulate atmospheric retrieval as a data-driven surrogate modeling problem, replacing iterative optimization with direct inference. Specifically, we learn a parametric mapping

$$x = F_\theta(y, a) \tag{2}$$

where F_θ is a neural network, y denotes multi-channel satellite observations, and a represents auxiliary variables such as geolocation, viewing geometry, and surface parameters. This formulation bypasses explicit radiative transfer simulation and enables efficient, end-to-end estimation of atmospheric profiles.

To capture both local and spatially coherent structures in the atmosphere, we design MiRS-AI with two complementary architectures. First, a pixel-wise ResNet model is used to emulate the traditional 1D-Var retrieval process. ResNet is a residual network (He et al., 2016) composed of three residual blocks, which maps ATMS observations to atmosphere profiles, effectively representing the inverse of the earlier work that maps atmosphere profiles to sensor observations (Liang et al., 2026). The skip connections in ResNet enable the model to learn residual mapping, facilitating improved gradient flow and faster convergence compared to direct mapping. The model takes as input 22-channel ATMS brightness temperatures along with auxiliary geophysical variables and outputs vertical profiles of temperature and water vapor for each pixel independently. This design provides a direct comparison to the operational MiRS framework and serves as a baseline for evaluating the benefits of data-driven retrieval.

Second, to overcome the limitations of independent-pixel assumptions, we introduce a spatially aware UNT model that operates on 96 x 96 image patches. The UNT architecture is inspired by UTNet (Gao et al., 2021), originally developed for medical image segmentation. The overall architecture of the UNT is illustrated in Figure 1. We incorporate Transformer module into both the encoder and decoder at each stage of the UNet, including the downsampling and upsampling paths. The UNet component (Ronneberger et al., 2015) captures local spatial patterns through multi-scale convolutional encoding and decoding, while the Transformer component (Vaswani et al., 2017) models long-range depen-

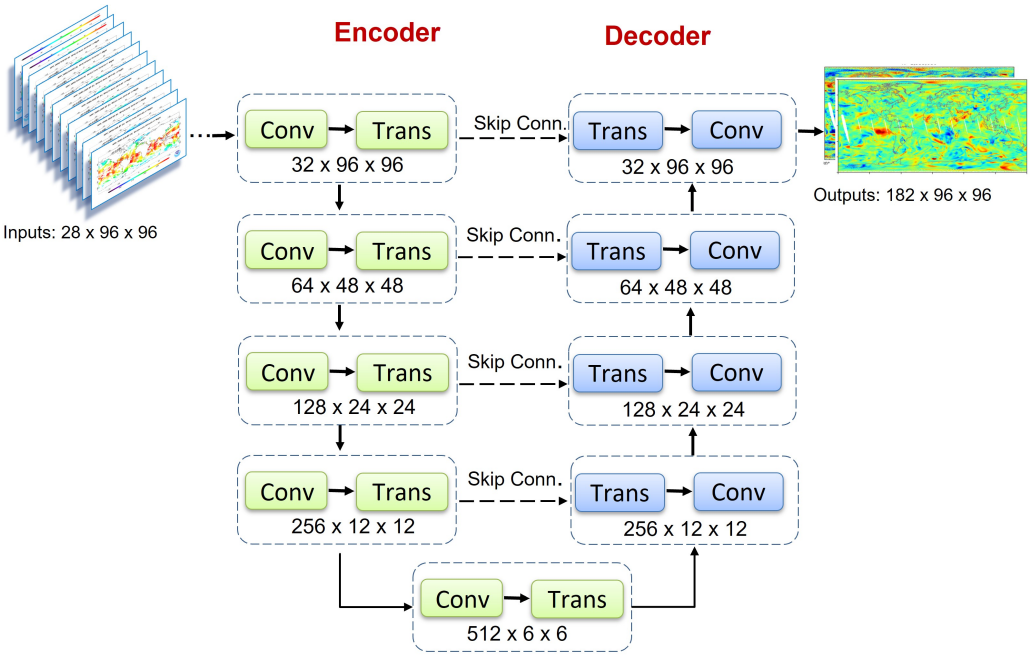

Figure 1. UNT architecture. (Conv: Convolution Block; Trans: Transformer block. Both blocks are defined as in Gao et al., 2021.).

dencies and provides global context within each input patch via self-attention mechanisms. By integrating both local and patch-level global information, the UNT model better represents spatial correlations in atmospheric fields, which is particularly important in dynamically complex environments such as tropical cyclones.

Both models are trained in a supervised manner using collocated satellite observations and analysis-derived atmospheric profiles. The training objective combines standard regression losses to apply for this task:

$$\mathcal{L} = \alpha \cdot MSE + \beta \cdot MAE \tag{3}$$

where *MSE* (mean squared error) captures overall error magnitude and *MAE* (mean absolute error) improves robustness to outliers. Physical consistency is learned implicitly from the large-scale training data.

This formulation allows MiRS-AI to serve as an efficient surrogate for variational retrieval, enabling fast and scalable atmospheric profile estimation while preserving key geophysical structures.

3. MiRS-AI Inputs and Labels

We construct a large-scale dataset by collocating satellite observations with ECMWF atmospheric analysis products to train and evaluate the proposed MiRS-AI framework. We use brightness temperature measurements from the sensor dataset record product (SDR) of the Advanced Technology Microwave Sounder (ATMS) onboard NOAA-20. For each observation, we extract all 22 microwave channels, which provide sensitivity to atmospheric temperature and water vapor across vertical layers. In addition to radiance measurements, we incorporate geophysical and surface-related variables to improve retrieval performance. These include: geolocation information (latitude and longitude); sensor viewing geometry (sensor zenith angle), and surface parameters such as surface temperature, surface pressure, and surface type extracted from ECMWF. These variables provide important ancillary information that supports the mapping from satellite measurements to temperature and water vapor profiles, similar to their role in radiative transfer–based retrieval.

Atmospheric state variables are obtained from quarter degree resolution and 91-layer ECMWF analysis data (Liang et al., 2026), which are collocated with ATMS observations at the pixel level. The target outputs consist of vertical profiles of atmospheric temperature and water vapor. These analysis products serve as high-quality reference labels for supervised learning.

To ensure coverage across diverse atmospheric conditions, we construct training datasets spanning multiple years and seasons. The ResNet model is trained using data from 2020–2021, sampling one day per month, while as the UNT model is trained using data from 2020–2022, sampling two days per month to support its higher data requirements for

spatial learning. A summary of the data inputs and labels are listed in Table 1. Overall, the dataset provides global coverage across all seasons, enabling the models to learn robust mappings under a wide range of atmospheric regimes.

4. Experiments and Results

4.1. Global Retrieval Performance

We first evaluate MiRS-AI on a globally distributed test datasets from October 27, 2025, which is nearly three years after the training period. We then evaluate the UNT model separately on additional dates, including January 4, 2023; March 30, 2023; June 30, 2023; and December 30, 2023, along with October 27, 2025, spanning multiple seasons and years (2023–2025). Retrieval performance is assessed by comparing estimated temperature and water vapor profiles against ECMWF analysis, using mean bias and standard deviation as primary metrics.

Figure 2 shows both the pixel-wise ResNet and the spatial UNT models consistently outperform the operational MiRS system across all test days. In particular, MiRS-AI generally achieves lower and smoother bias and reduced variance throughout the vertical profile, indicating improved accuracy and stability in the retrieval. Among the two models,

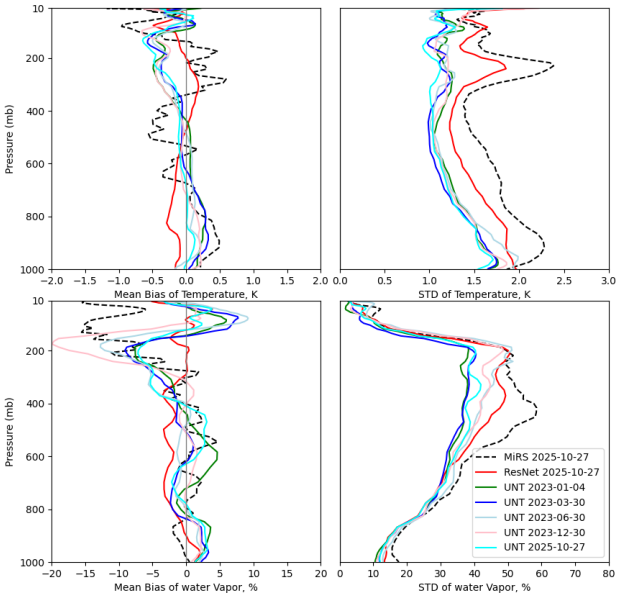

Figure 2. The mean biases (MiRS or MiRS-AI minus ECMWF analysis; left panels) and STDs (right panels) of UNT for 91-layer temperatures (upper panels) and water vapor (lower panels) in global for 5 days from 2023 to 2025 represent all 4 seasons. The red curve and dash black curve represent the results of ResNet and MiRS, respectively, on October 27, 2025.

the UNT model demonstrates the best overall performance, achieving the lowest standard deviation across most of different atmospheric layers and geographic regions. Importantly, these improvements are consistent across all seasons, suggesting strong generalization capability under diverse atmospheric conditions.

In the meantime, the UNT model consistently outperforms the pixel-wise ResNet, demonstrating that incorporating spatial context and multiscale dependencies significantly improves retrieval accuracy. This result highlights a key limitation of traditional 1D-Var systems and pixel-wise learning approaches, which neglect spatial correlations inherent in atmospheric fields.

4.2. Tropical Cyclone Case Study

Although UNT is designed as global-scale model trained on globally distributed data, this section focuses on evaluating its performance under dynamically complex conditions. To this end, we conduct case studies on tropical cyclones, including Hurricane ERIN, which traversed the Atlantic Ocean in August 2025, and Super Typhoon RAGASA, which developed over the Western Pacific Ocean in September 2025 (Figure 3). These scenarios present strong gradients, nonlinear dynamics, and multiscale interactions that are challenging for both traditional MiRS retrieval system and the MiRS-AI framework.

Figure 4 demonstrates that the UNT model achieves substantial improvements over MiRS across most vertical layers for both temperature and water vapor retrieval in the ERIN and RAGASA regions. These improvements underscore the importance of leveraging spatial context, highlighting that incorporating neighboring information is critical for accurate retrieval in dynamically complex environments.

Figure 5 presents the spatial distributions of temperature and water vapor at 300 mb, a level that typically exhibits strong signal for both variables within the vertical structure of tropical cyclones. Comparisons among MiRS, UNT, and ECMWF show that UNT produces smoother and more physically consistent retrieval fields, with reduced noise compared to MiRS, this indicates that the model effectively captures large-scale thermodynamic structures while maintaining spatial coherence. However, the spatial patterns also suggest that UNT may partially inherit the smoothing characteristics of the ECMWF analysis used as a reference. As a result, fine-scale thermodynamic structures, particularly within the inner-core regions of tropical cyclones, may be underrepresented. This implies that, in addition to relying on smoothed ECMWF analyses, future improvements to UNT may benefit from incorporating additional observation-driven constraints or higher-resolution references to better capture small-scale variability and extreme gradients in tropical cyclone environments.

Table 1. Input features and reference labels for ResNet and UNT.

Input Features			Reference Labels		
Name	**ResNet**	**UNT**	**Name**	**ResNet**	**UNT**
Sensor Zenith Angle	1	96 x 96	Temperature Profiles	91	96 x 96 x 91
Latitude	1	96 x 96			
Longitude	1	96 x 96			
Surface Type	1	96 x 96			
Surface Pressure	1	96 x 96	Water Vapor Profiles	91	96 x 96 x 91
Surface Temperature	1	96 x 96			
BTs	22	96 x 96 x 22			
Total	28	96 x 96 x 28	Total	182	96 x 96 x 182

Table 2. Efficiency comparison of MiRS and MiRS-AI ResNet and UNT for estimating temperature and water vapor profiles using one day (3.2 million) samples (CPU time in seconds).

3.2 million samples (CPU time, seconds)	MiRS	ResNet	UNT
T and WV estimate	6016	522	117

4.3. Computational Efficiency

A major limitation of variational retrieval systems such as MiRS is the high computational cost associated with iterative optimization and repeated evaluations of the radiative transfer model and its Jacobians. In contrast, MiRS-AI replaces this process with a single forward pass through a neural network, significantly reducing computational complexity.

As a result, MiRS-AI achieves orders-of-magnitude speedup compared to MiRS (Table 2), enabling near-real-time retrieval over large spatial domains. This improvement is particularly important for operational applications, where latency and scalability are critical. The efficiency gains also open new opportunities for integration with downstream systems. For example, MiRS-AI can serve as a fast surrogate within data assimilation frameworks (e.g., 3D/4D-Var), where repeated forward evaluations are required.

In addition, the framework can be extended to incorporate additional sensors, including microwave and hyperspectral observations, potentially using dimensionality reduction techniques such as PCA to manage high-dimensional inputs.

5. Conclusion

We present MiRS-AI, a deep learning framework for accurate and efficient atmospheric profile retrieval that replaces iterative variational optimization with direct inference. Experimental results demonstrate improved accuracy and significant computational efficiency compared to the operational MiRS system, with the spatial UNT model further enhancing performance in dynamically complex environments.

Despite these improvements, several limitations remain. Occasional outliers are observed in localized regions, likely associated with rare atmospheric conditions or distribution shifts. More fundamentally, the UNT model is trained as a global-scale model using ECMWF analysis as reference, which introduces an inherent smoothing bias. As a result, fine-scale thermodynamic structures—particularly within the inner-core regions of tropical cyclones—may be underrepresented. This reflects a broader limitation of data-driven surrogate retrievals trained on model-constrained targets, where the learned solution can inherit biases from the reference dataset.

Future work will focus on addressing these limitations through outlier detection and data balancing, and hybrid learning strategies that incorporate physical constraints and observation-driven consistency. In particular, integrating radiance-based supervision or higher-resolution references may help recover fine-scale structures while preserving large-scale coherence. Extensions to all-sky retrieval, multimodal sensor fusion, and the incorporation of MiRS-AI components into MiRS represent promising directions for advancing fast and physically consistent atmospheric retrieval systems.

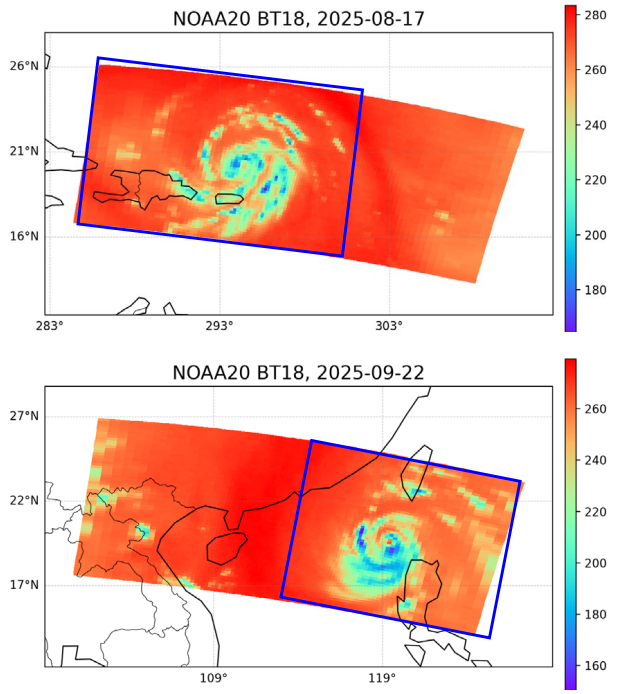

Figure 3. The mean biases (MiRS or MiRS-AI minus ECMWF analysis; left panels) and STDs (right panels) of UNT for 91-layer temperatures (upper panels) and water vapor (lower panels) in global for 5 days from 2023 to 2025 represent all 4 seasons. The red curve and dash black curve represent the results of ResNet and MiRS, respectively, on October 27, 2025.

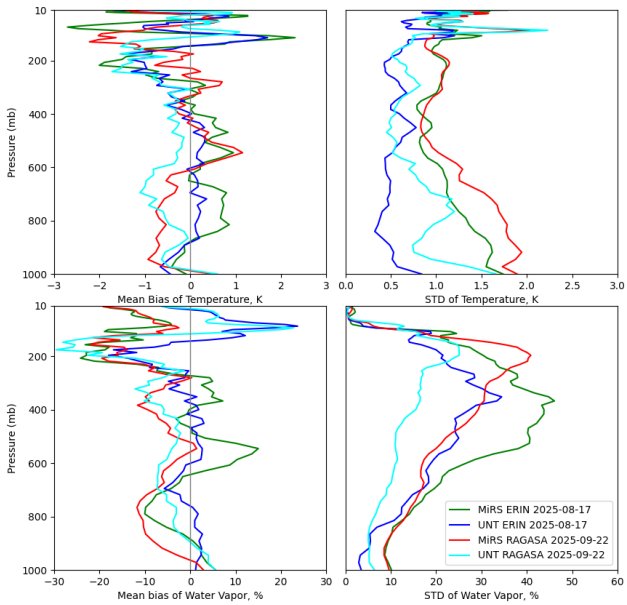

Figure 4. The mean biases (left) and STDs (right) of UNT and MiRS for 91-layer temperature (upper) and water vapor (lower) for the selected areas surrounding the tropical cyclone ERIN and RAGASA areas.

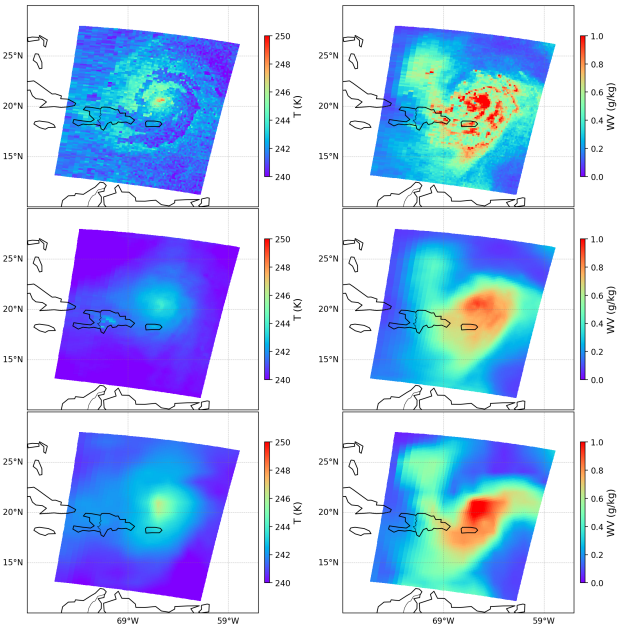

Figure 5. Comparison of temperature (left) and water vapor (right) at 300 mb for MiRS (top), UNT (middle), ECMWF (bottom) in the Hurricane ERIN region.

Acknowledgements

This research was supported by NOAA grant NA24NESX432C0001 (Cooperative Institute for Satellite Earth System Studies - CISESS) at the University of Maryland/ESSIC. The authors would like to thank the NOAA STAR CRTM Team for their valuable discussions and continued support. The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official NOAA or U.S. Government position, policy, or decision.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Boukabara, S. A., Garrett, K., Chen, W. C., Iturbide-Sanchez, F., Grassotti, C., Kongoli, C., Chen, R. Y., Liu, Q., Yan, H., Weng, F. Z., Ferraro, R., Kleespies, T. J., and Meng, H. Mirs: An all-weather 1dvar satellite data

assimilation and retrieval system. *IEEE Transactions on Geoscience and Remote Sensing*, 49(9):3249–3272, 2011.

Gao, Y., Zhou, M., and Metaxas, D. Utnet: A hybrid transformer architecture for medical image segmentation. *Medical Image Computing and Computer Assisted Intervention*, 2107.00781:61–71, 2021.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 770–778, Las Vegas, NV, 2016.

Lee, Y. K., Liu, Q., Grassotti, C., Liang, X., Liu, S., and Zhou, Y. Preliminary development and testing of an eps-sg microwave sounder proxy data generator using the noaa microwave integrated retrieval system. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:3151–3161, 2021.

Liang, X., Garrett, K., Liu, Q., Maddy, E. S., Ide, K., and Boukabara, S. A deep-learning-based microwave radiative transfer emulator for data assimilation and remote sensing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:8819–8833, 2022.

Liang, X., Liu, Q., and Grassotti, C. Refining clear-sky ai radiative transfer model forward predictions and jacobian accuracy for atms using a resnet and physical constraints. *IEEE Transactions on Geoscience and Remote Sensing*, 64:1–15, 2026.

Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *Med. Image Comput. Comput.-Assist. Intervent.–MICCAI*, pp. 234–241, Munich, Germany, 2015.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Wang, X., Parrish, D., Kleist, D., and Whitaker, J. Gsi 3dvar-based ensemble–variational hybrid data assimilation for ncep global forecast system: Single-resolution experiments. *Mon. Wea. Rev.*, 141:4098–4117, 2013.