



RACCOONN: Retrievals of Atmospheric Conditions Computed using Observations & Optimized Neural Networks

tremblaybenoit/RACCOONN
Benoit.Tremblay@ec.gc.ca

Benoit Tremblay¹, Joël Bédard¹, Mark Buehner¹

¹ Environment & Climate Change Canada, Science and Technology Branch, Meteorological Research Division, Dorval, Canada

1. INTRODUCTION: RETRIEVALS

Motivation: Satellite brightness temperatures have a significantly positive impact on atmospheric analyses and ensuing numerical weather predictions. However, much information content remains untapped due to certain limitations (e.g., clouds, surface-sensitive radiances), and only a subset of available channels is assimilated for numerical weather predictions.

Retrievals: Convert radiances into thermodynamic profiles that are more readily assimilated, but retrieval methods generally rely on strong assumptions.

Proposition: A deep learning framework that emphasizes a data-driven approach to retrievals from satellite radiances, with less dependence on the predefined assumptions of classical methodologies such as 1D-Var. The study aims to evaluate whether this shift in paradigm enables the assimilation of additional radiance (e.g., extra channels, cloudy scenes) in the form of retrieved thermodynamic profiles.

2. METHODOLOGY: NEURAL FIELD MODEL

Neural fields: Neural networks \vec{N} that represent a physical state \vec{x} as a continuous function of coordinates \vec{c} , i.e. high-dimensional curve-fitters and interpolators [1]:

$$\vec{x}(\vec{c}) \approx \vec{N}(\vec{c}).$$

\vec{N} is **mesh-independent**, treats time as a spatial coordinate (**no time marching**), and automatically computes **adjoints** and **spatial derivatives** with respect to \vec{c} [2]. However, \vec{N} is valid within the coordinate range spanned by \vec{c} during training (e.g., a given time window) and must likely be retrained outside of that domain.

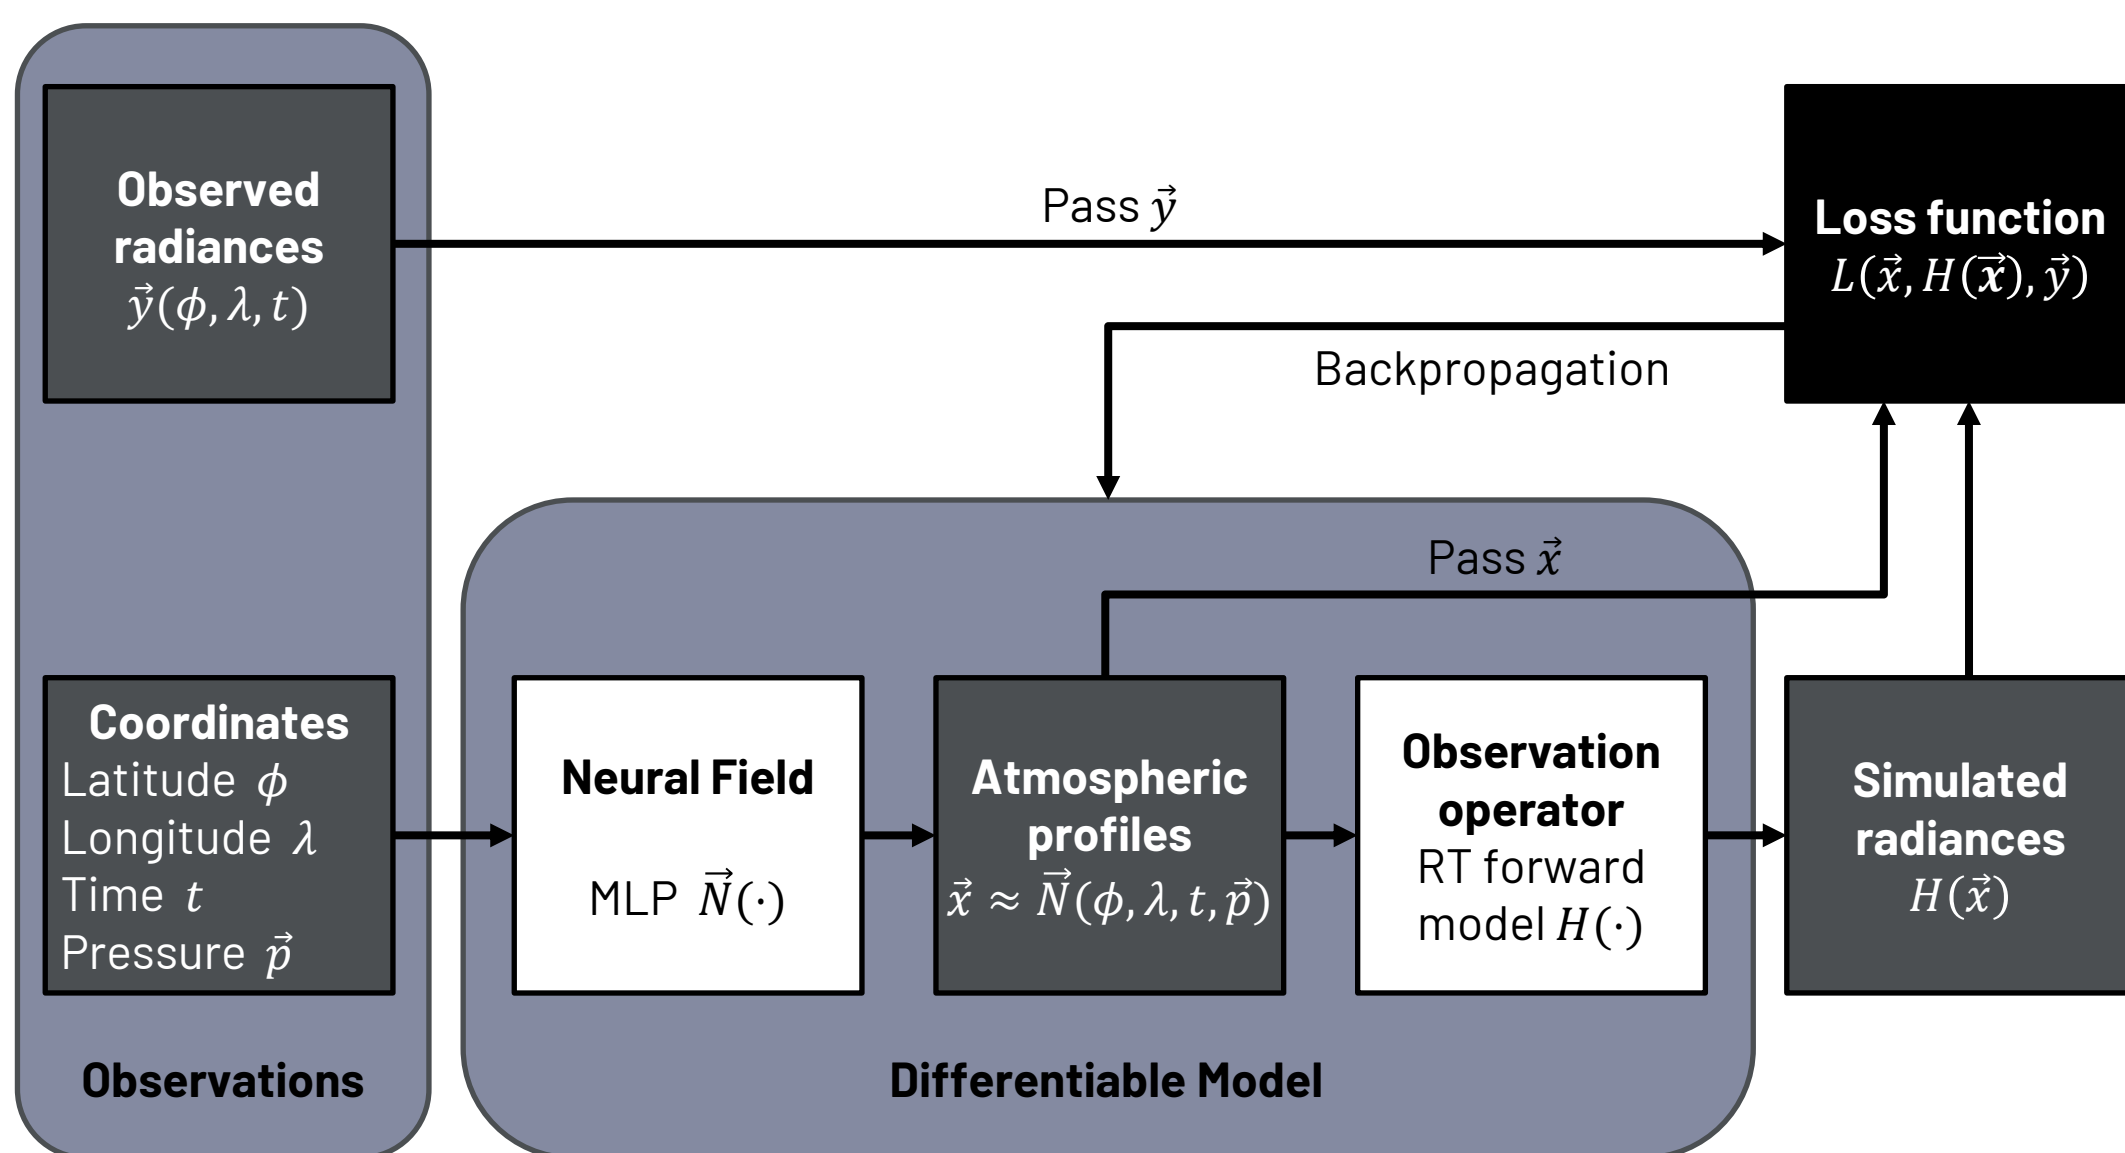


Fig. 1: RACCOONN schematic. Leverage a multi-layer perceptron (MLP) neural field \vec{N} and **observed radiances** $\vec{y}(\phi, \lambda, t)$ observed at time t , latitude ϕ , longitude λ to retrieve **atmospheric profiles** $\vec{x}(\phi, \lambda, t, \vec{p})$ over pressure levels \vec{p} . Note: **Observed radiances \vec{y} are not inputs to the network.** Instead, they appear directly in the loss function designed in a form akin to variational data assimilation cost functions:

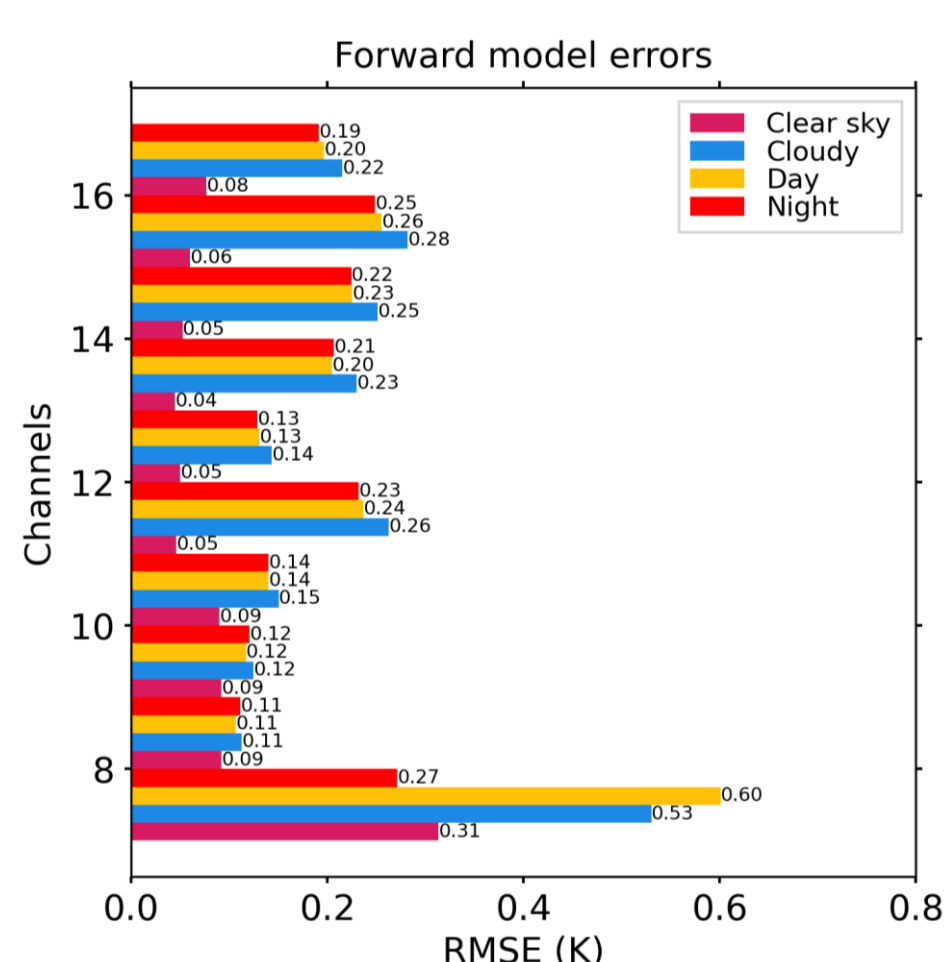
$$L(\vec{x}, H(\vec{x}), \vec{y}) = \alpha_{\text{obs}}(H(\vec{x}) - \vec{y})^T \mathbf{R}^{-1}(H(\vec{x}) - \vec{y}) + \alpha_{\text{model}}(\vec{x} - \vec{x}_b)^T \mathbf{B}^{-1}(\vec{x} - \vec{x}_b).$$

An observation loss term compares the reference brightness temperatures \vec{y} to the temperatures simulated from the neural-network-modeled profiles \vec{x} using a **fully differentiable radiative transfer (RT) forward model H** [3]. A model loss term regularizes the inferred profiles \vec{x} to a prior \vec{x}_b .

3. LOSS FUNCTION: RADIATIVE TRANSFER OPERATOR

Forward model: The probabilistic, fully-differentiable emulator of the Community Radiative Transfer Model (CRTM) from [3] maps atmospheric profiles $\vec{x}(\phi, \lambda, t, \vec{p})$, surface variables, and meta data to GOES/ABI radiances $\vec{y}(\phi, \lambda, t)$ (channels 7-16). It was re-implemented in PyTorch Lightning as an observation operator H to enable automatic adjoint computation [2].

Fig. 2: RMSE between the CRTM GOES/ABI radiance channels and the emulator inferences as a function of the sky conditions and time of day.



4. PROOF OF CONCEPT: SYNTHETIC DATA

Setup: Atmospheric profiles \vec{x}_t from a numerical simulation are used to simulate radiances $\vec{y} \approx H(\vec{x}_t)$ using Section 3's forward model H . The ground truth profiles \vec{x}_t are then reserved strictly for evaluation, while \vec{y} is fed to the loss function.

Scope: The experiment is restricted to a single timestep (static) and to clear sky conditions. The model infers air temperature and humidity and ozone mixing ratios. We assume a climatological prior \vec{x}_b , from which we also derive a \mathbf{B} matrix.

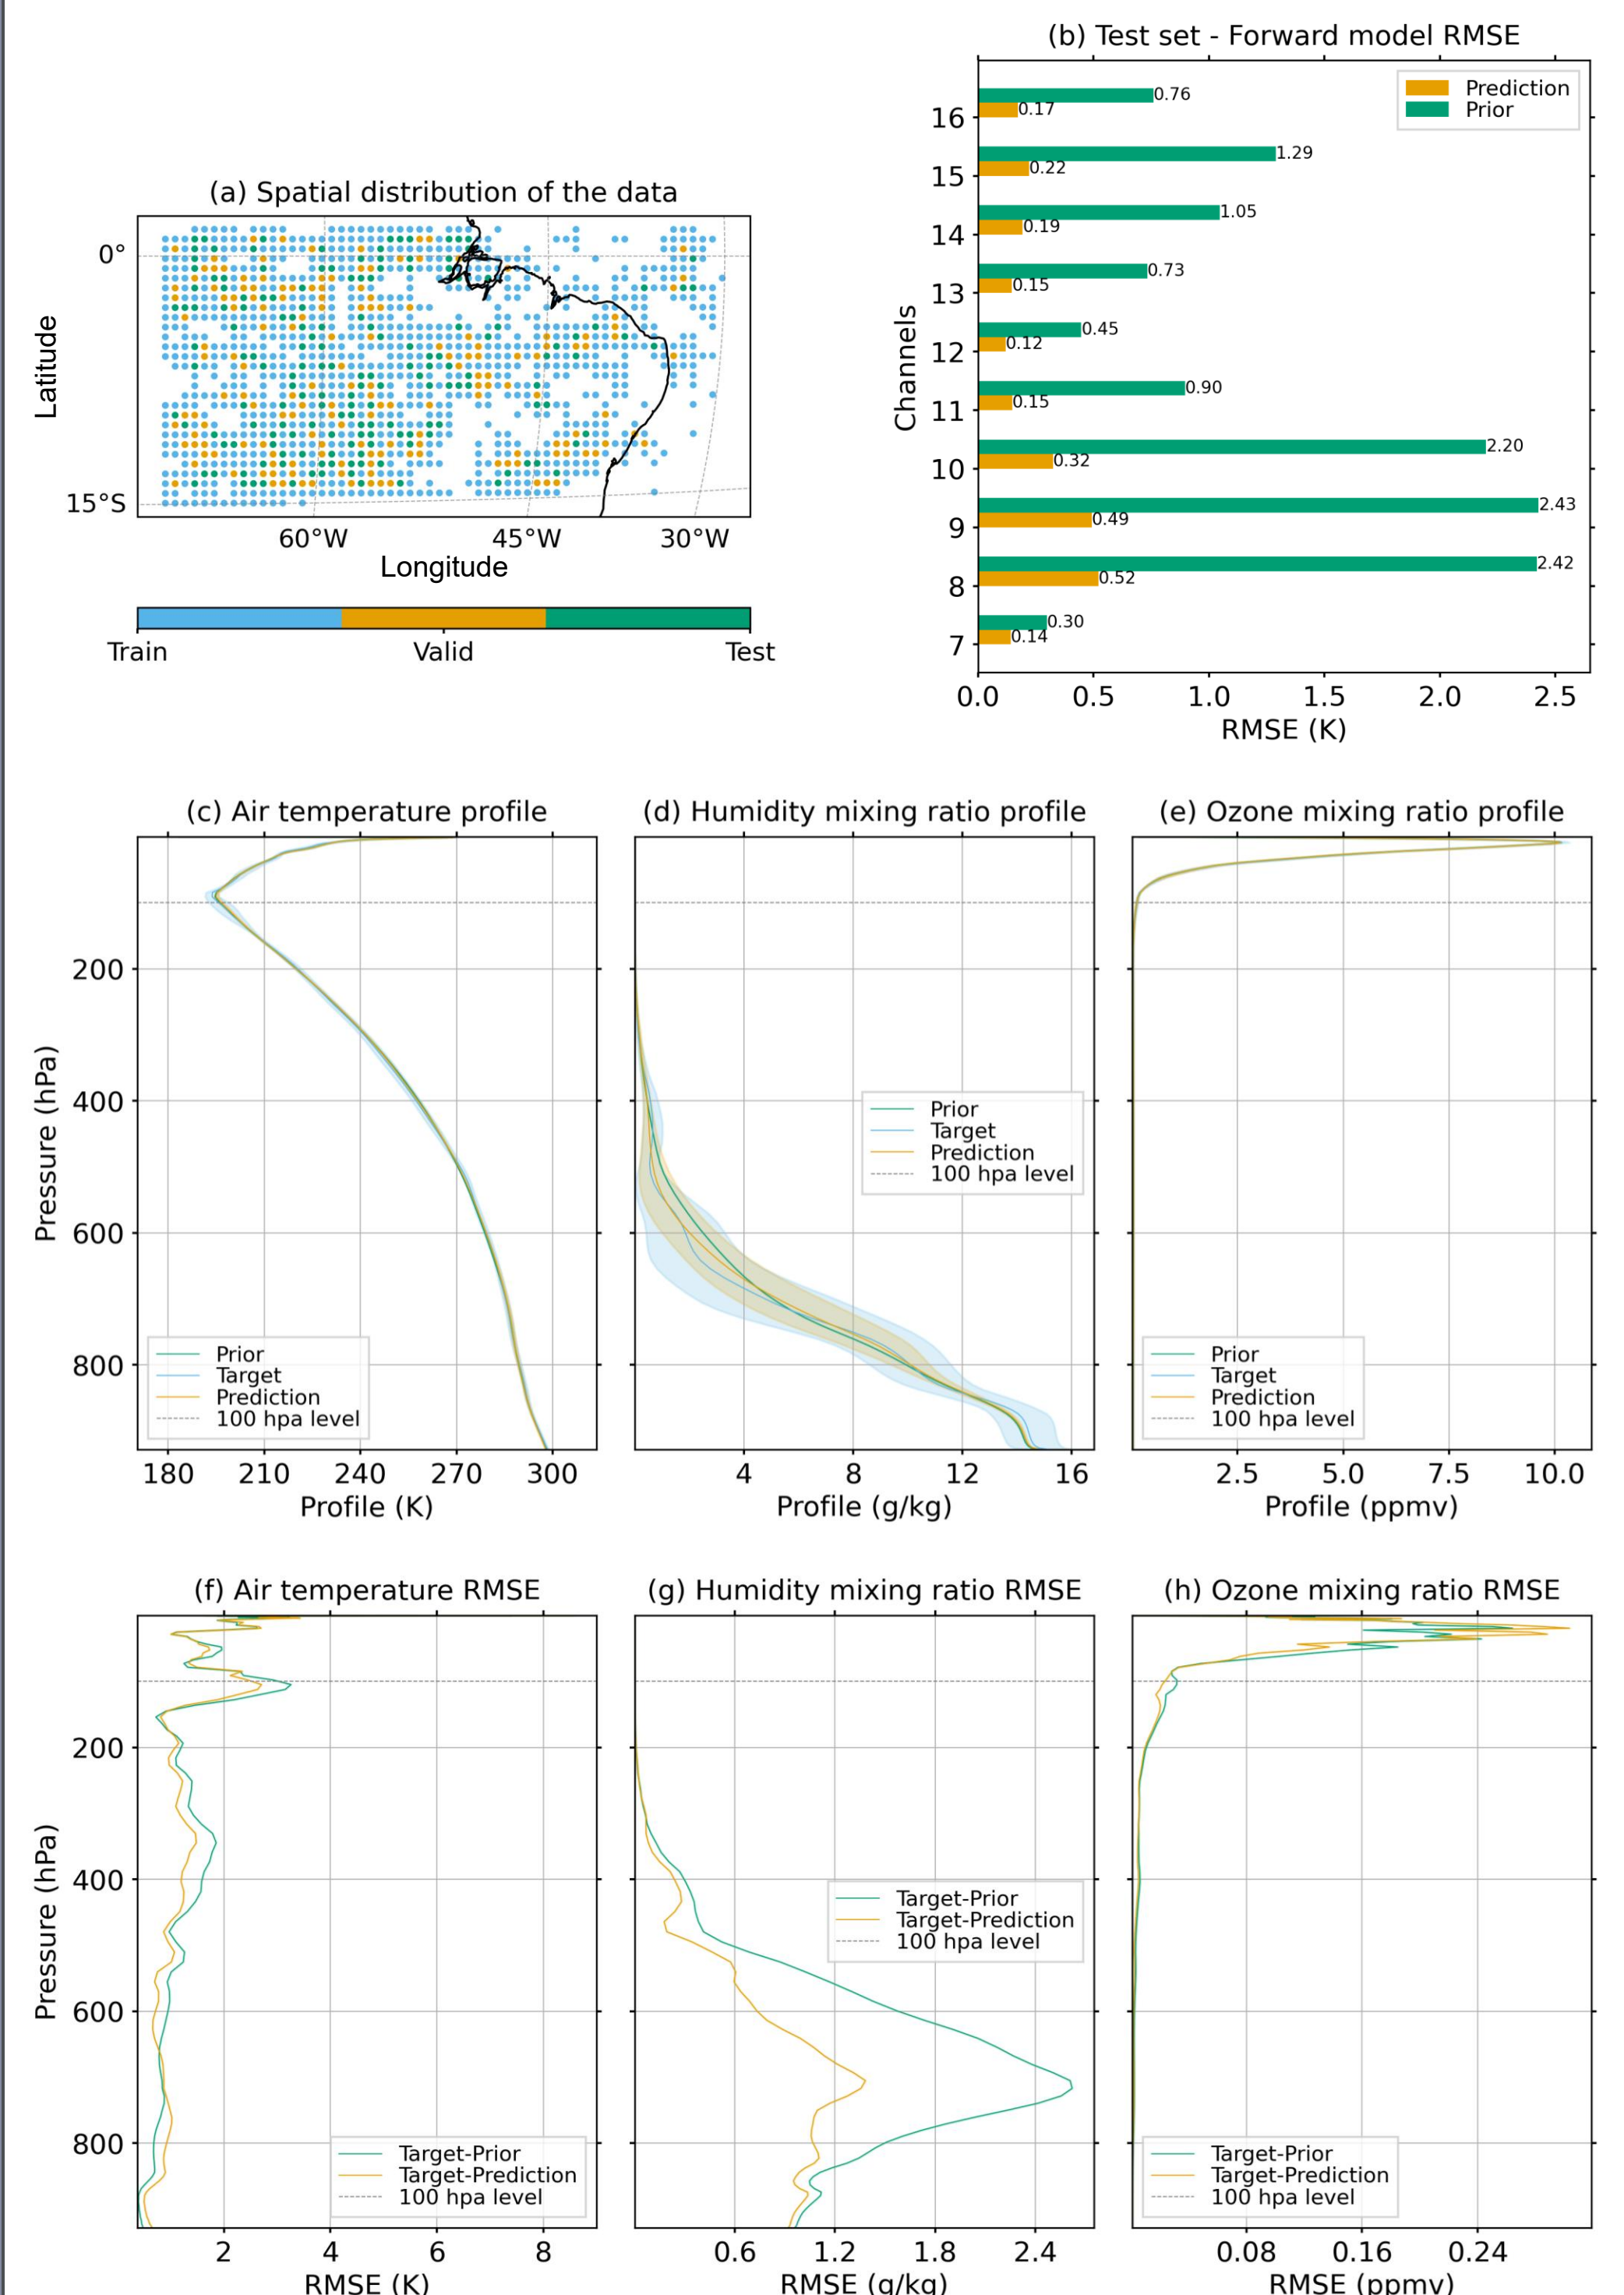


Fig. 3: Test set - (a) Scope and spatial distribution of the data. (b) RMSE between the (true) synthetic brightness temperatures and the ones forward-modeled from RACCOONN profiles. (c) through (h): Atmospheric profiles and their RMSE as a function of pressure level. The climatological prior is represented by the green curve. The mean and standard deviation of the (true) synthetic profiles are represented by the blue curve and blue shaded areas around it. Similarly, the mean and spread of the RACCOONN-inferred profiles are represented in orange.

5. CONCLUSIONS & FUTURE WORK

Summary: A neural field is used to retrieve thermodynamic profiles directly from observed radiances, providing a data-driven alternative to variational retrievals.

Short term goal: Assess whether the proposed data-driven approach can relax the assumptions of 1D-Var while matching or exceeding its performance.

Long-term goal: If successful, assimilate the neural field retrievals.

6. ACKNOWLEDGEMENTS & REFERENCES

The synthetic data and the source code for the forward model were kindly provided by Lucas Howard (CU Boulder), Aneesh C. Subramanian (CU Boulder), Gregory Thompson (JCSDA), Benjamin Johnson (JCSDA), and Thomas Auligné (JCSDA).

[1] Xie et al. (2022): <https://doi.org/10.48550/arXiv.2111.11426>.

[2] Jarolim et al. (2025): <https://doi.org/10.48550/arXiv.2502.13924>.

[3] Howard et al. (2025): <https://doi.org/10.48550/arXiv.2504.16192>.