
CuMoLoS-MAE: A Masked Autoencoder for Remote Sensing Data Reconstruction

Anurup Naskar¹ Nathanael Zhixin Wong¹ Sara Shamekh¹

¹ Courant Institute of Mathematical Sciences, New York University
{an4462,nw2648,ss18284}@nyu.edu

Abstract

Accurate atmospheric profiles from remote sensing instruments such as Doppler Lidar, Radar, and radiometers are frequently corrupted by low-SNR (Signal to Noise Ratio) gates, range folding, and spurious discontinuities. Traditional gap filling blurs fine-scale structures, whereas deep models lack confidence estimates. We present **CuMoLoS-MAE**, a Curriculum-Guided Monte Carlo Stochastic Ensemble Masked Autoencoder designed to (i) restore fine-scale features such as updraft and downdraft cores, shear lines, and small vortices, (ii) learn a data-driven prior over atmospheric fields, and (iii) quantify pixel-wise uncertainty. During training, CuMoLoS-MAE employs a mask-ratio curriculum that forces a ViT decoder to reconstruct from progressively sparser context. At inference, we approximate the posterior predictive by Monte Carlo over random mask realisations, evaluating the MAE multiple times and aggregating the outputs to obtain the posterior predictive mean reconstruction $\bar{X}$ together with a finely resolved per-pixel uncertainty map σ_X . Together with high-fidelity reconstruction, this novel deep learning-based workflow enables enhanced convection diagnostics, supports real-time data assimilation, and improves long-term climate reanalysis.

1 Introduction

Understanding climate change’s impacts on extreme events requires accurate and continuous measurements from remote sensing instruments such as Doppler-lidar, radar, or radiometers. However these data often suffer from missing or corrupted returns that must be filled in before they can be reliably used in downstream applications. Classical gap-filling procedures, such as sliding-window mean filters [1], smear out critical small-scale features like shear lines and updraft cores. Recent advances in deep-learning approaches (e.g. variational autoencoders [2]) can recover sharper structures but provide no information to gauge reconstruction uncertainties when these methods are used in downstream assimilation and alerting frameworks.

To tackle these shortcomings, we propose **CuMoLoS-MAE**, a Curriculum-Guided Monte Carlo Stochastic Ensemble Masked Autoencoder, a novel approach that comprises three core mechanisms: (1) curriculum masking [3] to stabilise training and encourage reconstruction from sparser context, (2) micro-patch based on masked autoencoder [4] to capture fine structure and mid-scale dynamics, and (3) Monte-Carlo [5] ensembles to produce per-pixel uncertainty maps alongside the final reconstructions. In this paper, we have used Doppler-lidar measurements from the Southern Great Plains (SGP) provided by the Atmospheric Radiation Measurement (ARM) [6] as a case study, and present our reconstructions of the vertical velocity field and their corresponding uncertainty quantification.




Figure 1: From Doppler lidar time–height arrays we form 64×64 images. During training we randomly hide a subset of patch tokens (the mask ratio increases from 0.5 to 0.7 over a curriculum), add positional encodings, pass the visible tokens through a ViT encoder, and reconstruct the full field with a lightweight decoder. The loss is computed only on the hidden pixels. At test time, for each unseen patch we draw 50 independent random masks and run the same pipeline to produce multiple reconstructions. We then average these reconstructions to obtain a single best denoised estimate of the field and use the pixel-wise spread of the ensemble as an uncertainty map.

2 Data

We have used vertical velocity data from a Doppler lidar at the ARM SGP site, stored as time–height arrays in NetCDF format with 320 range gates at 30 m spacing, where each range gate represents the mean return from a fixed vertical layer. Focusing on the lowest 64 gates (1.92 km), we extract non-overlapping 64×64 patches as inputs for training and evaluation. The model is trained on data from June 1–9, 2011 and evaluated on an unseen day: June 15, 2011.

3 Methodology

3.1 CuMoLoS–MAE Architecture

Inspired by masked autoencoding and stochastic ensembling methods for uncertainty quantification, the CuMoLoS-MAE framework is comprised of two stages as shown in Figure 1:

- **Micro–patchified MAE with curriculum:** Each 64×64 Doppler–velocity slice is tokenized into 2×2 micro–patches. A 12–layer ViT [7] encoder processes visible tokens and a 4–layer decoder reconstructs the field. During training, the mask ratio r starts at 50% for 5 epochs, then cosine–ramps to 70% by epoch 30 and is held thereafter; optimization uses masked–MSE on the hidden pixels.

- **Monte Carlo ensembling:** At inference time, a new random mask is generated for each forward pass and the masking, encoding, and decoding process is repeated $N = 50$ times. The resulting ensemble of predictions is then aggregated to compute the mean and pixel-wise standard deviation:

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N \hat{X}^{(i)}, \quad \sigma_X = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{X}^{(i)} - \bar{X})^2} \quad (1)$$

yielding a high-fidelity denoised reconstruction $\bar{X}$ and per-pixel uncertainty estimate σ_X .

3.2 Training Procedure

The data are preprocessed by applying an SNR filter (intensity ≥ 0.005) and clamping valid velocities to the range $[-5, 5] \text{ m s}^{-1}$. From the resulting fields, non-overlapping 64×64 patches are extracted. Training is conducted over 500 epochs on a single NVIDIA A100 GPU using the AdamW optimiser [8] (base learning rate $1.5 \times 10^{-4} \cdot \frac{32}{256}$, weight decay 0.05) with a batch size of 32. A cosine learning rate schedule with warmup aligned to the mask-ratio curriculum period is employed.

3.3 Uncertainty map

At inference, $N = 50$ independent random masks are drawn per image, each masked input is passed through the decoder, and the resulting ensemble is aggregated (Eq. (1)) to obtain the posterior mean $\bar{X}$ and the pixel-wise standard deviation σ_X . The value of N is chosen to balance ensemble stability with computational cost.

4 Results

On the 1028 held-out test set, CuMoLoS-MAE surpasses the 8×8 mean filter, Noise2Void U-Net [9], Noise2Void Denoising CNN (DnCNN) [10], and a Convolutional Variational Autoencoder (CVAE [11])(Table 1). It also attains the highest low-frequency spectral fidelity (93.75%). Spectral integrity compares temporal Power Spectral Density of the denoised output $P_{\text{den}}(f)$ and raw signal $P_{\text{raw}}(f)$ via:

$$\varepsilon_{\log}(f) = \frac{\log_{10} P_{\text{den}}(f) - \log_{10} P_{\text{raw}}(f)}{\log_{10} P_{\text{raw}}(f)}. \quad (2)$$

Fidelity is the fraction of bins with $|\varepsilon_{\log}(f)| \leq 0.5$ for $f \leq 0.01$ Hz. CuMoLoS-MAE satisfies this in 93.75% of low-frequency bins, preserving storm-scale energy. Gate 0 is the principal exception, likely owing to its location at the base of the atmospheric boundary layer as shown in Figure 2. The DnCNN (Noise2Void) baseline yields the best FID (0.12) despite worse reconstruction and lower spectral fidelity, likely due to perceptual similarity in low-variance regions.

Table 1: Reconstruction performance and low-frequency spectral fidelity on 1028 held-out radar patches. Up arrow ($\uparrow$) indicates higher is better; down arrow ($\downarrow$) indicates lower is better.

Method	PSNR (dB) $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	FID $\downarrow$	Spectral Fidelity $\uparrow$
8×8 Mean-Filter	23.41	0.4950	0.5186	5.13	91.67%
CVAE	26.70	0.4190	0.4036	3.28	80.21%
DnCNN (Noise2Void)	23.09	0.6466	0.6232	0.12	36.46%
U-Net (Noise2Void)	27.70	0.7016	0.2581	0.44	49.48%
CuMoLoS-MAE (ours)	29.45	0.7857	0.1854	1.87	93.75%

4.1 Uncertainty Quantification and Climate Utility

The reliability of the uncertainty estimates is assessed by computing the pixel-wise Pearson correlation between the Monte Carlo standard deviation map σ_X and the absolute reconstruction error. Across 1028 test images, strong alignment is observed, with a mean per-patch correlation of $r = 0.961 \pm 0.037$, a global correlation of $r = 0.961$, and a Spearman rank correlation of $\rho = 0.926$. As illustrated in Figure 2, these high correlations indicate that σ_X closely tracks the true reconstruction error, i.e., the model reliably predicts where its outputs are likely to be accurate and where they are not. These pixels are split into ten bins by σ_X and the Mean Absolute Error (MAE) in each is computed. The MAE rises monotonically from 0.02845 to 0.99939 across σ -deciles (a $35.1 \times$ gap), and the top 1, 5, 10, and 20 percent of pixels by σ capture 10.1, 30.6, 43.4, and 59.4 percent of total |error|, confirming the utility of σ_X for error triage in downstream assimilation and warning systems.

4.2 Effect of Temporal Context on Reconstruction and Spectral Fidelity

Table 2: Reconstruction quality and spectral fidelity across different temporal window sizes. Up arrow ($\uparrow$) indicates higher is better; down arrow ($\downarrow$) indicates lower is better.

Window (time x range gate)	PSNR (dB) $\uparrow$	SSIM $\uparrow$	MSE $\downarrow$	FID $\downarrow$	Spectral Fidelity $\uparrow$
64×64	29.45	0.7857	0.1854	1.87	93.75%
128×64	30.11	0.7697	0.2253	3.73	87.50%
256×64	28.55	0.6103	0.3205	5.50	38.02%

From Table 2, reconstruction quality and low-frequency spectral fidelity decline as the window increases from 64×64 to 128×64 and then 256×64 , indicating a predominantly local denoising regime in which 64×64 provides sufficient context. The slight PSNR bump at 128×64 arises from a few near-constant, ultra-low-error intervals overweighted by the log average (MSE remains nearly unchanged), which is not statistically significant for any meaningful inference. At 256×64 , token count and masked area rise without added capacity, increasing errors across metrics.




Figure 2: Visual diagnostics for one sample. Left: log-PSD at selected range gates—reconstruction matches the original (small gaps) except at Gate 0 near the boundary-layer base. Centre-left: original Doppler-lidar vertical velocity. Centre-right: **CuMoLoS-MAE** reconstruction. Right: per-pixel uncertainty map σ_X .

4.3 Effect of Mask-Ratio Curriculum on Training Efficiency

Curriculum masking accelerates convergence while preserving perceptual quality. Reconstruction loss falls below 0.20 by epoch 198 (vs. 224 for fixed; 13 min faster) and reaches 0.189 by epoch 286 (vs. 333). As shown in Table 3, metrics are comparable, so we conclude that curriculum masking improves training efficiency by $\sim 10\%$ despite marginally better final scores with fixed masking.

Table 3: Effect of mask-ratio curriculum on reconstruction, FID, and spectral fidelity (1028 patches). Up arrow ($\uparrow$) indicates higher is better; down arrow ($\downarrow$) indicates lower is better.

Configuration	PSNR (dB) $\uparrow$	MSE $\downarrow$	SSIM $\uparrow$	FID $\downarrow$	Spectral Fidelity $\uparrow$
Without curriculum	29.45	0.1854	0.7857	1.87	93.75%
With mask-ratio curriculum	28.90	0.2106	0.7868	1.88	93.23%

5 Conclusion and Future Work

We have introduced CuMoLoS-MAE, a curriculum-guided masked autoencoder with Monte Carlo ensembling, which delivers state-of-the-art reconstructions together with per-pixel uncertainty. By recovering fine-scale features lost to noise or masking and flagging low-confidence regions, CuMoLoS-MAE improves detection of atmospheric coherent structures. This capability is valuable for Earth observation and climate modelling, as more accurate reconstructed fields improve (1) our understanding of extreme events and their evolution under global warming, (2) data assimilation in numerical weather prediction systems, and (3) process-level analysis within atmospheric sciences. Uncertainty maps further enable reconstitution by weighting observations according to confidence.

Specifically, we plan to use this model to create datasets of vertical profiles to study the shallow-to-deep transition of convection. Reconstructing denoised vertical profiles and their associated uncertainties will help with (a) updraft detection, (b) better characterize their width and kinetic energy, and (c) improve predictors of the transition timing and likelihood. We will evaluate the generalizability of our model across lidar systems provided by ARM (e.g. new datasets from Bankhead National Forest) and assess real-time deployment for operational assimilation. We also plan to scale training from days to months and years to capture climatological variability, support long-term reanalysis, and test robustness under changing environmental conditions.

References

- [1] Lesti G, Spiegel S. A Sliding Window Filter for Time Series Streams. In: IOTSTREAMING@PKDD/ECML; 2017. .
- [2] Pinheiro Cinelli L, Araújo Marins M, Barros da Silva EA, Lima Netto S. Variational autoencoder. In: Variational methods for machine learning with applications to deep networks. Springer; 2021. p. 111-49.
- [3] Madan N, Ristea NC, Nasrollahi K, Moeslund TB, Ionescu RT. Cl-mae: Curriculum-learned masked autoencoders. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision; 2024. p. 2492-502.
- [4] He K, Chen X, Xie S, Li Y, Dollár P, Girshick R. Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition; 2022. p. 16000-9.
- [5] Harrison RL. Introduction to monte carlo simulation. In: AIP conference proceedings. vol. 1204; 2010. p. 17.
- [6] Sisterson D, Peppler R, Cress T, Lamb P, Turner D. The ARM southern great plains (SGP) site. Meteorological Monographs. 2016;57:6-1.
- [7] Han K, Wang Y, Chen H, Chen X, Guo J, Liu Z, et al. A survey on vision transformer. IEEE transactions on pattern analysis and machine intelligence. 2022;45(1):87-110.
- [8] Llugsi R, El Yacoubi S, Fontaine A, Lupera P. Comparison between Adam, AdaMax and Adam W optimizers to implement a Weather Forecast based on Neural Networks for the Andean city of Quito. In: 2021 IEEE Fifth Ecuador Technical Chapters Meeting (ETCM). IEEE; 2021. p. 1-6.
- [9] Song TA, Yang F, Dutta J. Noise2Void: unsupervised denoising of PET images. Physics in Medicine & Biology. 2021;66(21):214002.
- [10] Zhao Y, Li Y, Dong X, Yang B. Low-frequency noise suppression method based on improved DnCNN in desert seismic data. IEEE Geoscience and Remote Sensing Letters. 2018;16(5):811-5.
- [11] Bao J, Chen D, Wen F, Li H, Hua G. CVAE-GAN: fine-grained image generation through asymmetric training. In: Proceedings of the IEEE international conference on computer vision; 2017. p. 2745-54.