

Global Particulate Matter Forecasting Using Lightweight, Region-Specific Deep Learning Models

ICLR 2026
ML4RS Workshop

Ansh Kushwaha¹ & Kaushik Gopalan^{1,2}

¹Centre for Interdisciplinary AI (CAI), FLAME University ²School of Computing and Data Sciences, FLAME University

Approach

We train **30 independent U-Net models** (10 spatial tiles \times 3 PM species) to forecast PM_1 , $PM_{2.5}$, and PM_{10} at $\Delta t \in \{6, 12, 24\}$ h. Each model is specialised for one geographic region, capturing region-specific emission dynamics that unified global models fail to resolve.

Architecture & Training

Data: CAMS reanalysis 2021–2024, $0.4^\circ \times 0.4^\circ$ (450 \times 900 grid), 4 synoptic times/day. Ten surface-level inputs. 2024 held out for testing.

Nested-domain tiling: 10 overlapping 256 \times 256 tiles; central 192 \times 192 is the forecast region. 32-cell border provides spatial context. Predictions stitched into a global 369 \times 900 forecast.

U-Net: 5-level encoder–decoder (16 \rightarrow 32 \rightarrow 64 \rightarrow 128 \rightarrow 256 ch.), DoubleConv blocks, MaxPool2D, bilinear upsampling, skip connections, Tanh output.

Size: 4.3 M params per model (\approx 120 M total) vs. Aurora \approx 1 B params.

Training: AdamW ($\eta=10^{-3}$, $\lambda=10^{-4}$), cosine LR annealing, early stopping (patience 16); \approx 60 epochs. Composite loss: LogCosh + Huber + SSIM + MAE.

Conclusions

- ★ 30 lightweight U-Nets (\sim 120 M params) match or outperform Aurora (\sim 1 B params) for PM_1 & $PM_{2.5}$ at a fraction of the computational cost.
- ★ Nested-domain tiling is essential — removing it raises PM_1 CRMSE by up to 50% at 6 h.
- ★ U-Net achieves superior SSIM across all species and all forecast horizons.
- ★ Aurora’s PM_{10} advantage at 12–24 h reflects synoptic-scale coarse-particle transport.

Results — CRMSE / RMSE ($\mu\text{g m}^{-3}$) & SSIM

		6 h Forecast			12 h Forecast			24 h Forecast		
Metric	Model	PM_1	$PM_{2.5}$	PM_{10}	PM_1	$PM_{2.5}$	PM_{10}	PM_1	$PM_{2.5}$	PM_{10}
CRMSE	Persistence	3.4	6.0	30.2	4.2	8.8	40.2	4.3	7.8	31.5
	Aurora	4.7	7.4	36.4	3.0	6.5	28.0	3.4	6.1	25.9
	UNet (ours)	3.2	5.2	23.3	3.2	5.7	32.6	3.3	6.5	30.9
	<i>UNet w/o Nesting</i>	4.8	5.8	24.7	4.5	8.4	35.4	4.8	7.7	28.4
RMSE	Persistence	3.4	6.0	30.2	4.2	8.8	40.2	4.3	7.8	31.5
	Aurora	8.3	7.5	39.8	3.1	7.1	28.5	3.4	7.1	26.7
	UNet (ours)	3.2	5.2	23.3	3.2	5.8	32.6	3.3	6.5	30.9
	<i>UNet w/o Nesting</i>	4.8	5.8	24.7	4.5	8.4	35.4	4.9	7.8	28.5
SSIM	Persistence	0.95	0.94	0.98	0.94	0.94	0.97	0.93	0.90	0.96
	Aurora	0.70	0.93	0.96	0.97	0.97	0.98	0.94	0.89	0.96
	UNet (ours)	0.98	0.98	0.99	0.98	0.97	0.99	0.98	0.97	0.99
	<i>UNet w/o Nesting</i>	0.97	0.97	0.99	0.97	0.97	0.99	0.97	0.96	0.99

Bold = best per column (ties bolded for both). †Ablation without nested spatial context. Aurora not fine-tuned for PM.