

NeurIPS 2026 · Main Track

Wavelet Flow Matching for Multi-Scale Physics Emulation

LEAP
LEARNING THE EARTH WITH ARTIFICIAL INTELLIGENCE AND PHYSICS


Gabriele Accarino
Columbia/LEAP


Juan Nathaniel
Columbia/LEAP


Carla Roesch
Uni. of Edinburgh


Pierre Gentine
Columbia/LEAP


Sara Shamekh
NYU–Courant


D. Watson-Parris
UC San Diego/Scripps


Viviana Acquaviva
CUNY/LDEO



COLUMBIA | NYU | UC IRVINE | U. MINNESOTA | UC SAN DIEGO | CUNY | TEACHERS COLLEGE | NCAR | NASA | CarbonPlan

Emulating Chaotic Physical Dynamics

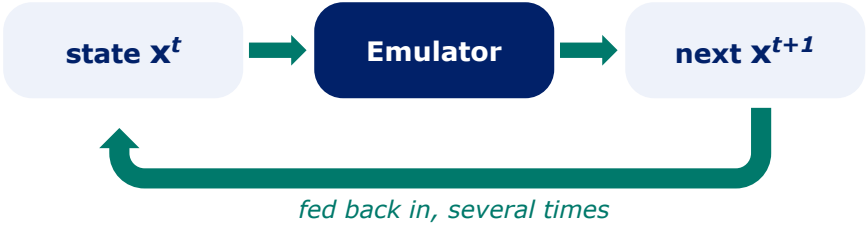

PDE: Turbulent Radiative Layer 2D

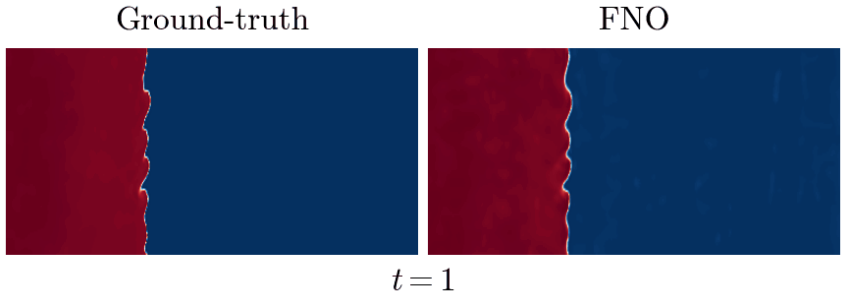

Deterministic emulators

U-Net, neural operators (e.g., FNO)

- ✓ Fast: one forward pass per step
- ✓ Accurate at short lead times
- ✗ Predict the mean future → blurry
- ✗ Errors compound → unstable

Generative, pixel space

diffusion / flow matching

- ✓ Sample the full distribution → sharp
- ✓ Ensembles → uncertainty
- ✗ Many network calls per step
- ✗ All at full resolution → slow

Generative, latent space

pre-trained autoencoder + generator

- ✓ Compact latent space → cheaper
- ✗ Extra stage: pre-train autoencoder
- ✗ Lossy: can drop fine-scale physics
- ✗ Accuracy capped by the decoder

Can we achieve sharp, stable autoregressive rollouts at low computational cost, *without learning a separate compression?*

NSF LEAP

Idea: Perform Generative Emulation in Wavelet Space

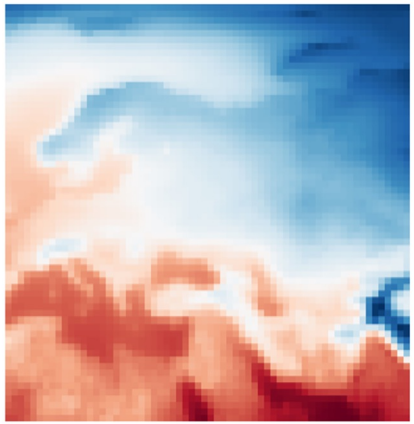

field x^t

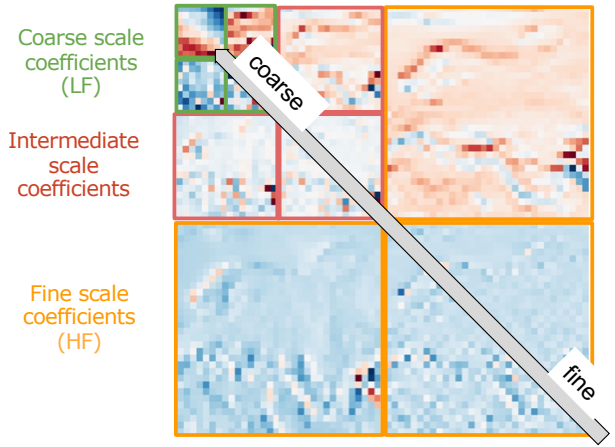

3-scale Discrete Wavelet Transform
$w^t = W(x^t)$

Usual: generate in pixel space

$$p(x^{t+1} \mid x^t)$$

Ours: generate in wavelet space

$$p(w^{t+1} \mid w^t) \rightarrow x^{t+1} = \mathcal{W}^{-1}(w^{t+1})$$

next coefficients *exact inverse, applied once*

Training-free

Prescribed basis, not learned: no autoencoder to pre-train

Lossless

Exact inverse transform: no information is thrown away

Local in space & scale

Where different frequencies occur and at what scales (Fourier modes are global)

Wavelet Flow Matching (WFM)

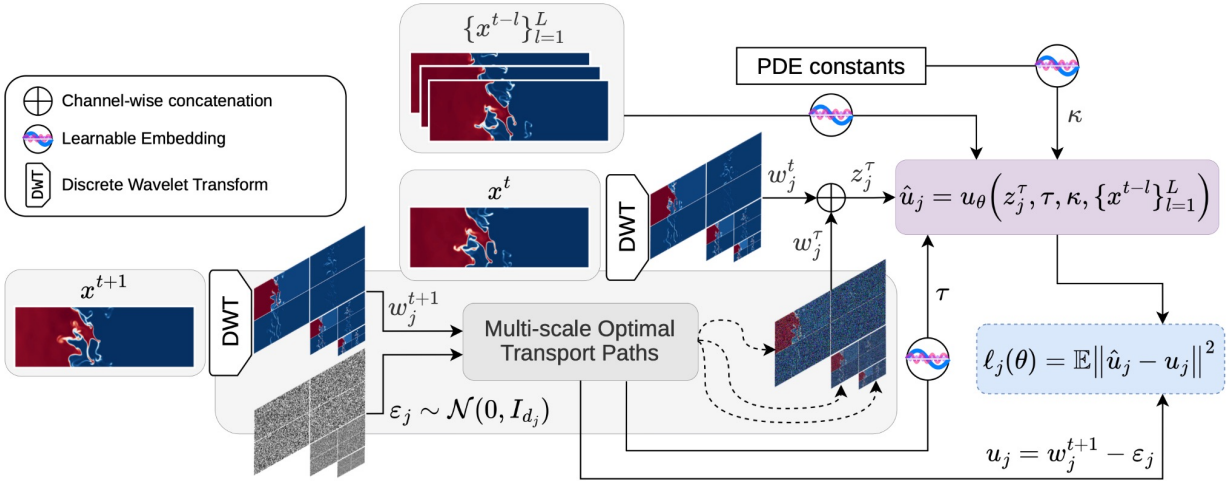

1 Decompose

The next state is split by a DWT into J scales; Gaussian noise is drawn at each scale.

2 Learn all scales jointly

One U-Net: each scale enters at its matching resolution and gets its own velocity head.

3 Sample & roll out

Integrate the flow (50 Euler steps), apply one inverse DWT, feed back.

Accarino et al., NeurIPS 2026 (Main Track)

NSF LEAP

Results: Autoregressive Rollouts

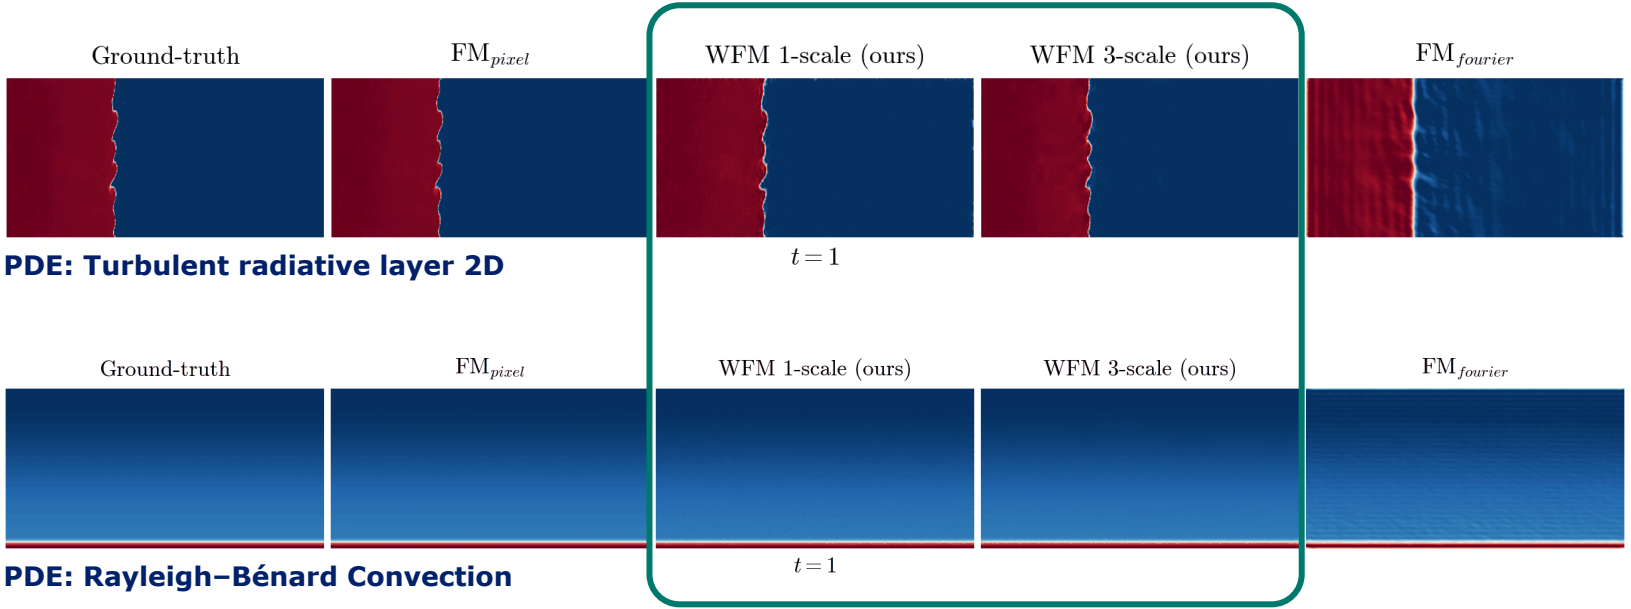

WFM keeps fine-scale structure sharp over long rollouts, matching pixel-space quality at lower cost, **while a global Fourier basis introduces artifacts.**

Results: WFM vs Deterministic and Generative Baselines

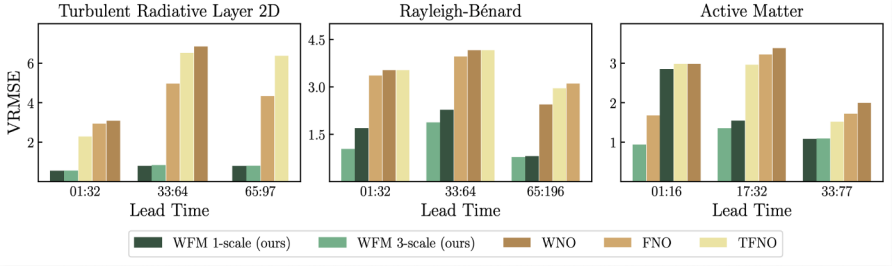

1 **vs. deterministic** - VRMSE (lower is better)

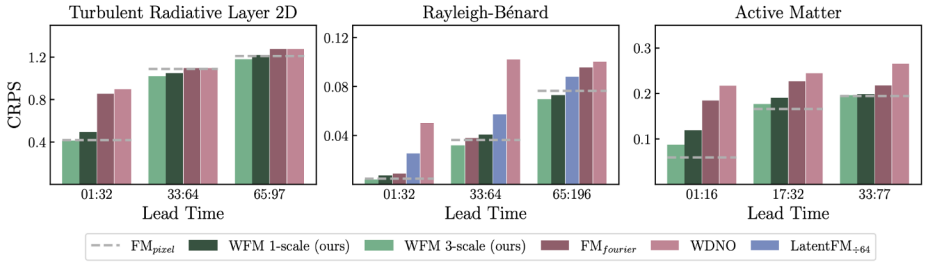

2 **vs. generative** - CRPS (lower is better)

3 **vs. generative** - spectral coherence RMSE (lower is better)

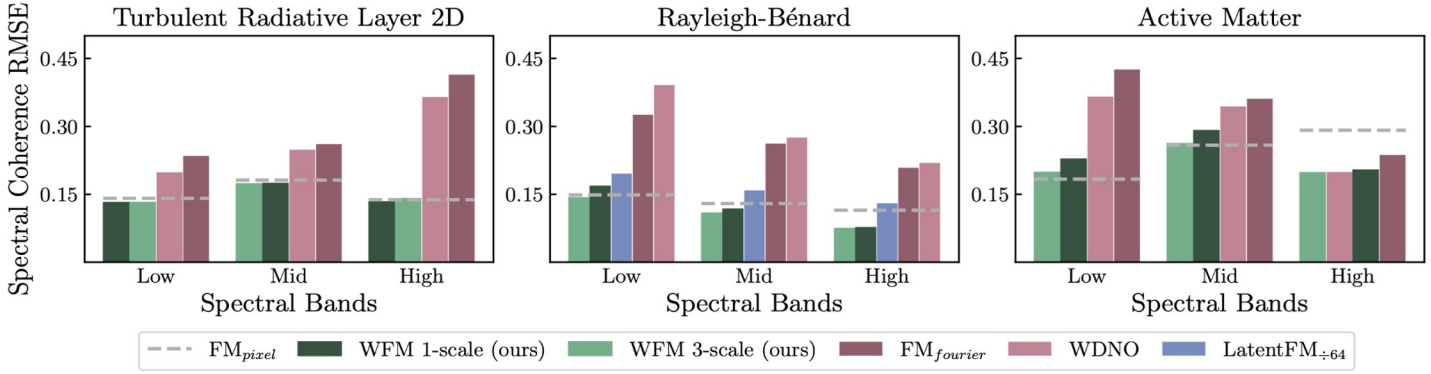

Accarino et al., NeurIPS 2026 (Main Track)

Takeaways

Wavelets are an effective, training-free space for generative emulation

Joint multi-scale flows give sharp, spectrally faithful, stable rollouts

Better accuracy–cost trade-off: no autoencoder, up to 1.5x faster than pixel space

Webpage

Project overview
+ Code
+ Pre-trained weights

Data

TheWell

ga2673@columbia.edu
https://gabrieleaccarino.github.io

📍 Wed, Dec 9, 2026, 12:30 PM – 2:30 PM CET Paris Poster Hall



LEAP

LEARNING THE EARTH WITH ARTIFICIAL INTELLIGENCE & PHYSICS



COLUMBIA | NYU | UC IRVINE | U. MINNESOTA | UC SAN DIEGO | CUNY | TEACHERS COLLEGE | NCAR | NASA | CarbonPlan