PODiff: Latent Diffusion in Proper Orthogonal Decomposition Space for Scientific Super-Resolution

Efficient, interpretable, and well-calibrated probabilistic super-resolution of scientific fields from low-resolution inputs

Onkar Jadhav, Tim French, Matthew Rayson, Nicole L. Jones

[1] School of Earth and Oceans, UWA Oceans Institute, The University of Western Australia

[2] School of Physics, Mathematics and Computing, The University of Western Australia

Paper (arXiv) — arXiv: 2605.03399

Code (GitHub) — OnkaraJadhav/PODiff

Interactive Demo — HF Demo: SST

THE UNIVERSITY OF WESTERN AUSTRALIA

ICML International Conference On Machine Learning

TL;DR

- **Diffusion in POD space** variance-ordered, orthogonal, interpretable latent space
- **Efficient & scalable:** much lower compute and memory than pixel-space diffusion
- **Well-calibrated uncertainty:** high empirical coverage and spatially meaningful uncertainty
- **State-of-the-art results:** strong accuracy with much lower cost

Motivation & Core Idea

■ Scientific fields, e.g., SST are high-dimensional and expensive to model directly in pixel space.

⏱ Pixel-space diffusion is accurate but often slow and memory-intensive.

💡 PODiff moves diffusion into a compact, variance-ordered POD latent space for efficient probabilistic super-resolution.

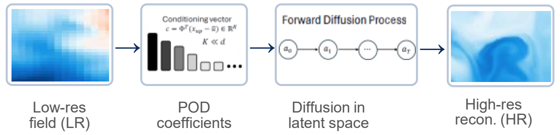

Key idea: denoise only the dominant POD coefficients, then reconstruct the full field.

PODiff Framework: Conditional Diffusion in POD Space

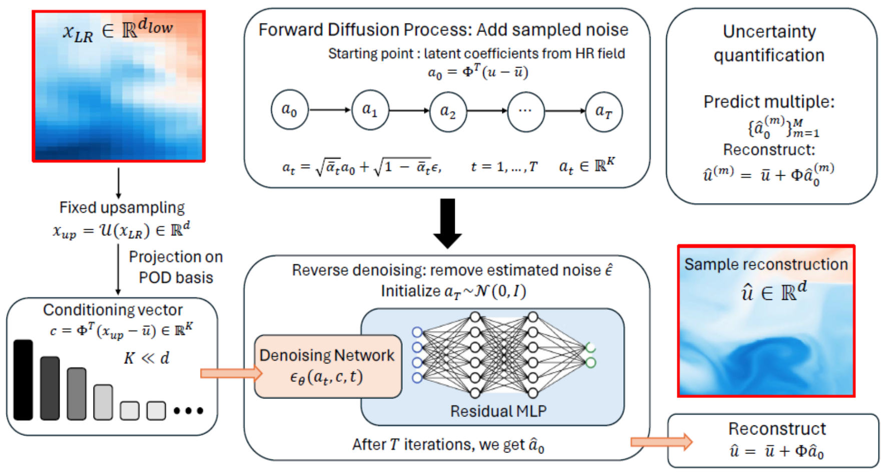

POD Modes Capture Multi-Scale Ocean Variability

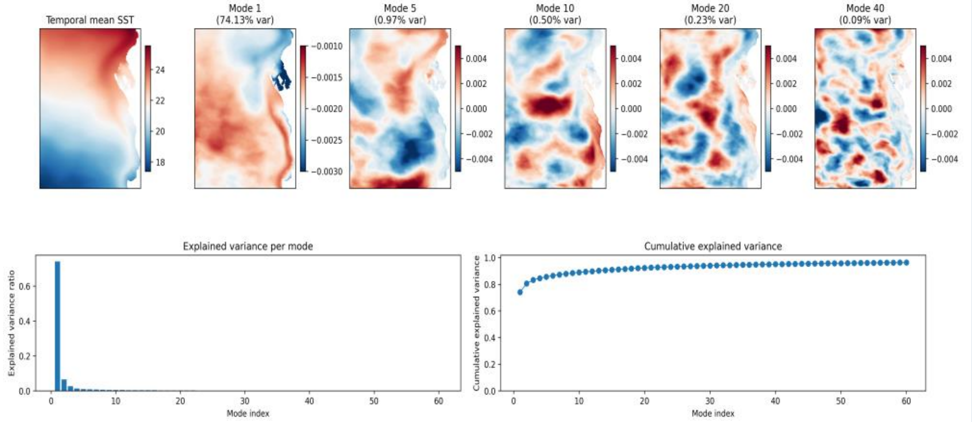

70%+ of total variance captured by the first POD mode alone!

SST Downscaling over Western Australian Coast

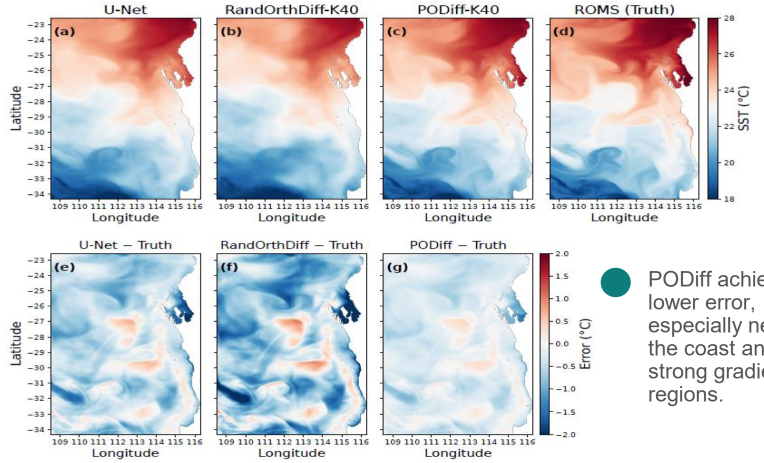

PODiff achieves lower error, especially near the coast and in strong gradient regions.

Well-Calibrated & Spatially Structured Uncertainty

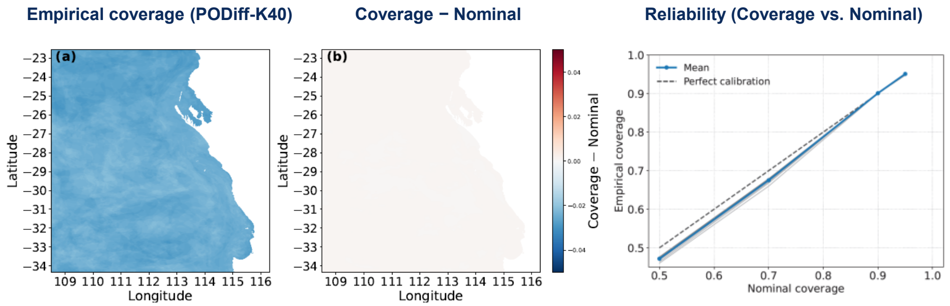

PODiff achieves high empirical coverage across all confidence levels with minimal bias.

Quantitative Results (SST Downscaling)

Reconstruction Error Metrics (All days)

MODEL	RMSE	MAE	EXTREME RMSE	EXTREME MAE
PODiff-K40 (ours)	**0.3923**	**0.2976**	**0.4836**	**0.3537**
PIXELDIFF	0.4118	0.3158	0.4899	0.3600
VAE-LDM	0.4011	0.3005	0.4889	0.3591
U-Net	0.6788	0.5141	0.8366	0.6109
RandOrthDiff	0.9987	0.7577	1.2309	0.9003

Empirical Coverage

Nominal Level	PODiff-K40	MC Dropout U-Net	PixelDiff
50%	0.4717	0.4111	0.4658
70%	0.6849	0.6508	0.6799
90%	0.9009	0.8871	0.9010
95%	0.9571	0.9401	0.9551

Computational Cost

Method	Params	Peak GPU Mem	Training Time	Inference Time per Sample
PODiff (K=40) (ours)	**0.20M**	**1.4 GB**	**3.8 h**	**0.08 s**
U-Net	33M	8.8 GB	8.2 h	0.05 s
PixelDiff	33M	12.5 GB	48 h	1.24 s
VAE-LDM	39M	13.1 GB	28 h	0.88 s

⚡ **PODiff is 30–100× faster** at inference and uses 2–5× less GPU memory than pixel-space diffusion!

Key Takeaway

By performing diffusion in a variance-ordered POD space, PODiff achieves state-of-the-art reconstruction accuracy with well-calibrated, spatially structured uncertainty at a fraction of the computational cost.

Applications

- Climate & ocean model downscaling
- Weather & extreme event analysis
- Data assimilation & digital twins
- Other PDE-governed spatial systems

Acknowledgments

We acknowledge funding from the State of Western Australia for the project “Advancing predictions of WA Marine Heatwaves and impacts on marine ecosystems”