

Topology-Inspired Fine-Tuning for Wind-Field Super-Resolution

Arjun Dadhwal, Brian Summa

Tulane University

Super-resolution: 2 km wind from 10 km input

~100 km

~10 km: SR input

~2 km: SR target

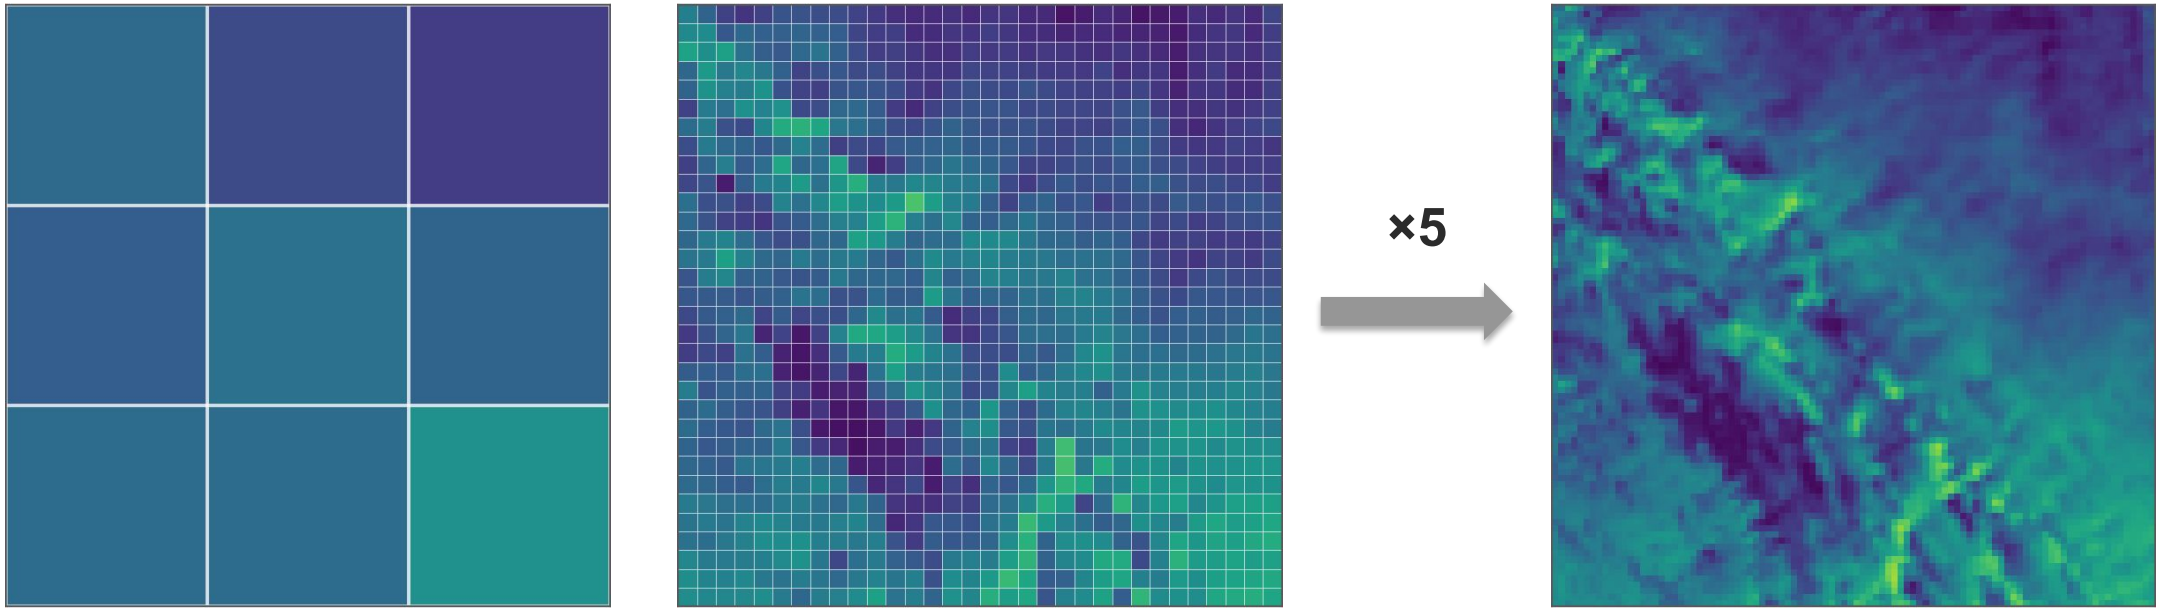

Wind-resource assessment needs ~2 km → **learn to downscale**

One WIND Toolkit field, block-averaged; not a climate-model prediction

Different metrics favor different reconstructions

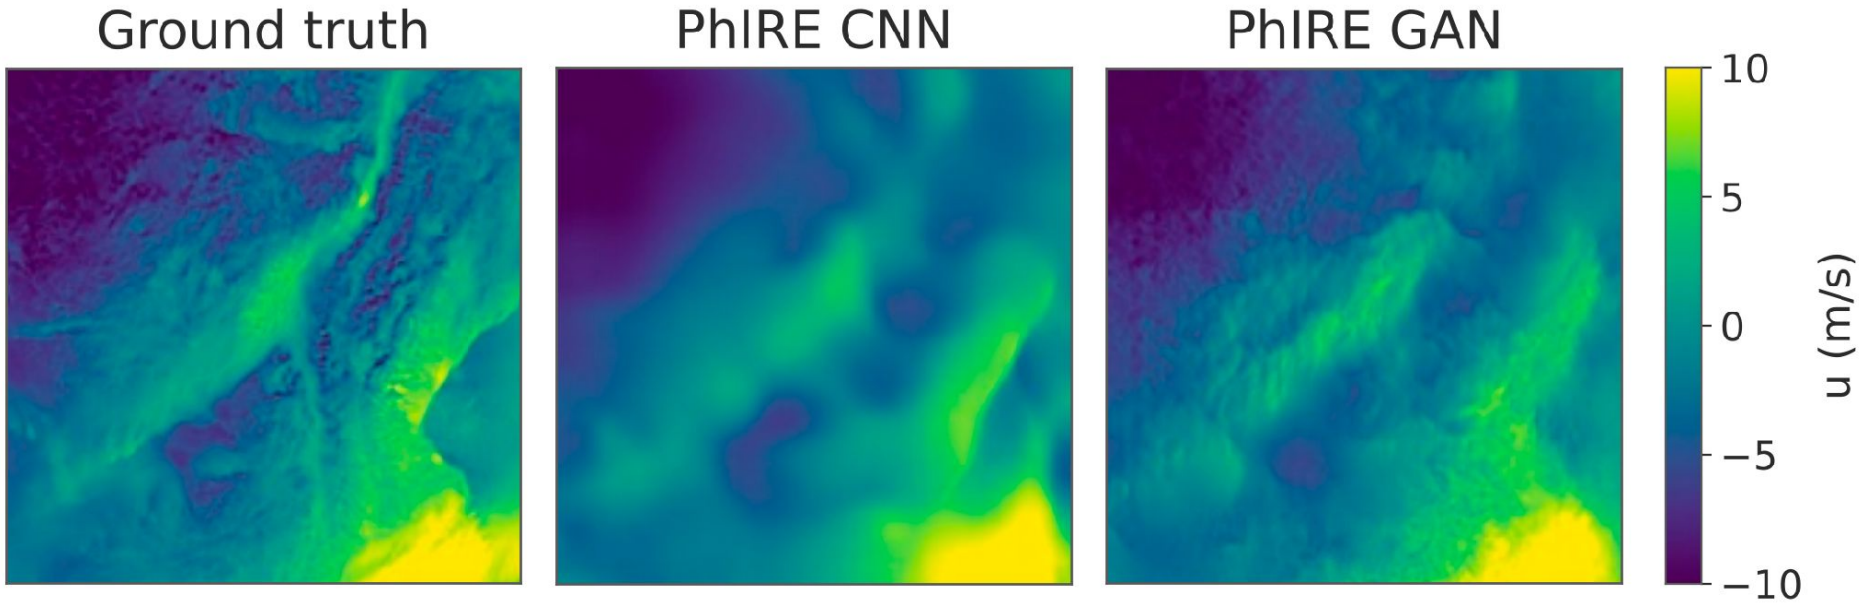

My 168-field benchmark PSNR: **CNN wins 168/168** PD $W_{2,2}$: **GAN wins 158/168**

Can we preserve persistent structure while keeping fidelity?

Fine-tuning with fixed GT persistence pairs

Pretrained CNN
same starting checkpoint

Control: standard fine-tuning

Ours: same + three terms below

Same evaluation
168 evaluation fields

Critical values

Match values at fixed GT birth and death points

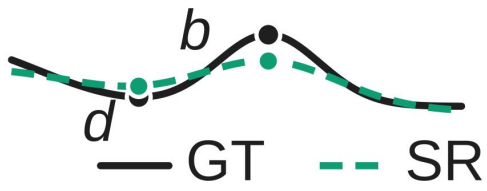

Persistence

Match each fixed GT pair's contrast |d − b|



Gradient

Match local wind-speed changes



GT pairs precomputed once with TTK · no diagram computed during training · PD and MT distances used only for evaluation

PD distances drop 25–31%; PSNR improves

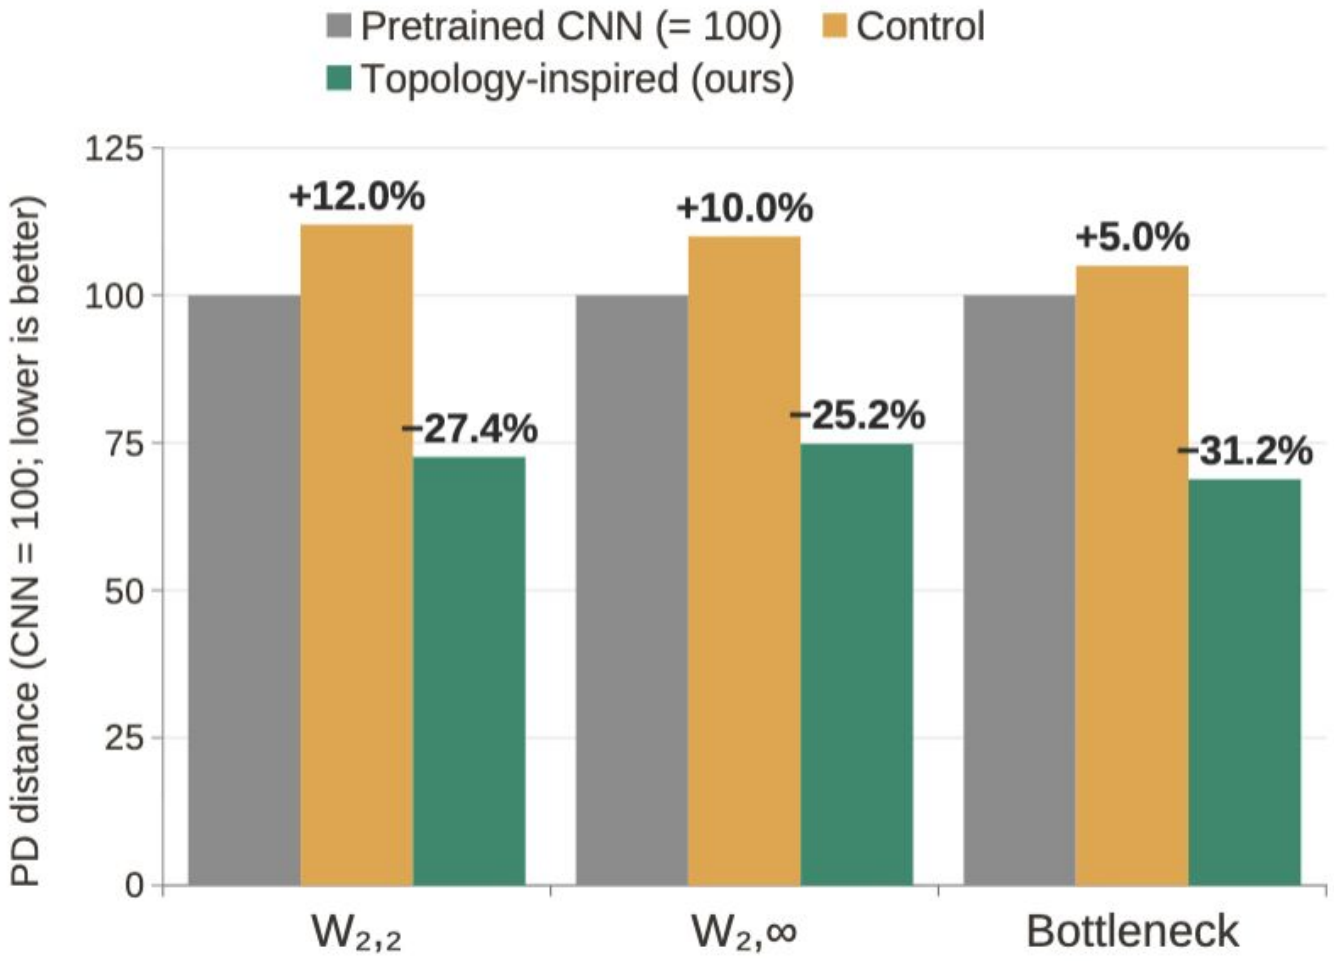

150/168

fields with all three PD distances lower than the pretrained CNN

+1.30 dB

PSNR vs. pretrained CNN (control: +2.60 dB, but worse PD)

61%

of those 150 also improve in merge-tree distance

Thank you!

Come see the poster for:

- Full qualitative comparison with the actual persistence diagrams
- All PD and merge-tree results
- Wind power, gradient and connectivity metrics

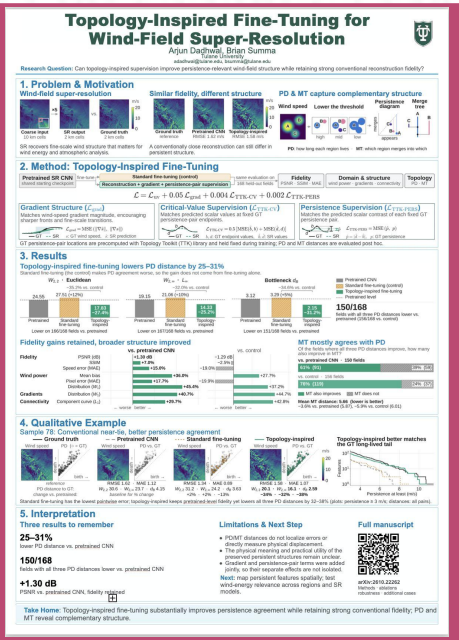

VIS2026

Full manuscript

ablations and more cases

arXiv:2610.22262

adadhwal@tulane.edu