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

Arjun Dadhwal, Brian Summa
Tulane University
adadhwal@tulane.edu, bsumma@tulane.edu

Research Question: Can topology-inspired supervision improve persistence-relevant wind-field structure while retaining strong conventional reconstruction fidelity?

1. Problem & Motivation

Wind-field super-resolution

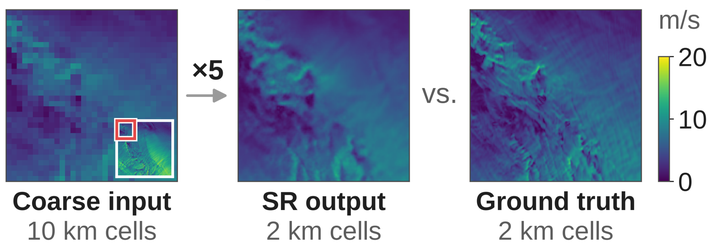

SR recovers fine-scale wind structure that matters for wind energy and atmospheric analysis.

Similar fidelity, different structure

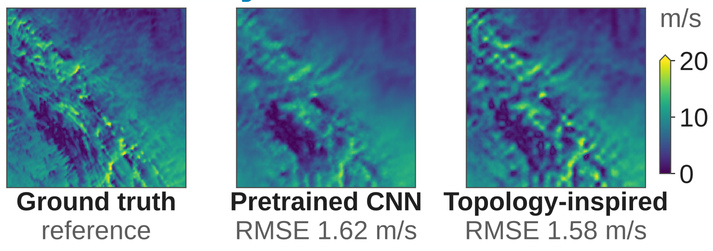

A conventionally close reconstruction can still differ in persistent structure.

PD & MT capture complementary structure

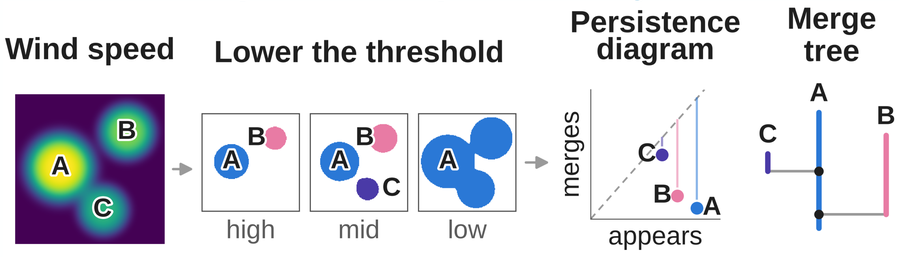

PD: how long each region lives · **MT**: which region merges into which

2. Method: Topology-Inspired Fine-Tuning

Pretrained SR CNN (shared starting checkpoint) → fine-tune → **Standard fine-tuning (control)** / **Reconstruction + gradient + persistence-pair supervision** → same evaluation on 168 held-out fields → **Fidelity** (PSNR · SSIM · MAE) · **Domain & structure** (wind power · gradients · connectivity) · **Topology** (PD · MT)

$$\mathcal{L} = \mathcal{L}_{\text{uv}} + 0.05\,\mathcal{L}_{\text{grad}} + 0.004\,\mathcal{L}_{\text{TTK-CV}} + 0.002\,\mathcal{L}_{\text{TTK-PERS}}$$

Gradient Structure ($\mathcal{L}_{\text{grad}}$)

Matches wind-speed gradient magnitude, encouraging sharper fronts and fine-scale transitions.



$$\mathcal{L}_{\text{grad}} = \text{MSE}\left(\|\nabla \hat{s}\|,\ \|\nabla s\|\right)$$

s: GT wind speed, $\hat{s}$: SR prediction

Critical-Value Supervision ($\mathcal{L}_{\text{TTK-CV}}$)

Matches predicted scalar values at fixed GT persistence-pair endpoints.

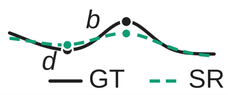

$$\mathcal{L}_{\text{TTK-CV}} = 0.5\left[\text{MSE}(\hat{b}, b) + \text{MSE}(\hat{d}, d)\right]$$

b, d: GT endpoint values, $\hat{b}, \hat{d}$: SR values

Persistence Supervision ($\mathcal{L}_{\text{TTK-PERS}}$)

Matches the predicted scalar contrast of each fixed GT persistence pair.

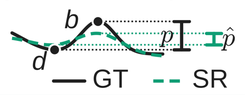

$$\mathcal{L}_{\text{TTK-PERS}} = \text{MSE}\left(\hat{p},\ p\right)$$

$\hat{p} = |\hat{d} - \hat{b}|$, p: GT persistence

GT persistence-pair locations are precomputed with Topology Toolkit (TTK) library and held fixed during training; PD and MT distances are evaluated post hoc.

3. Results

Topology-inspired fine-tuning lowers PD distance by 25–31%

Standard fine-tuning (the control) makes PD agreement worse, so the gain does not come from fine-tuning alone.

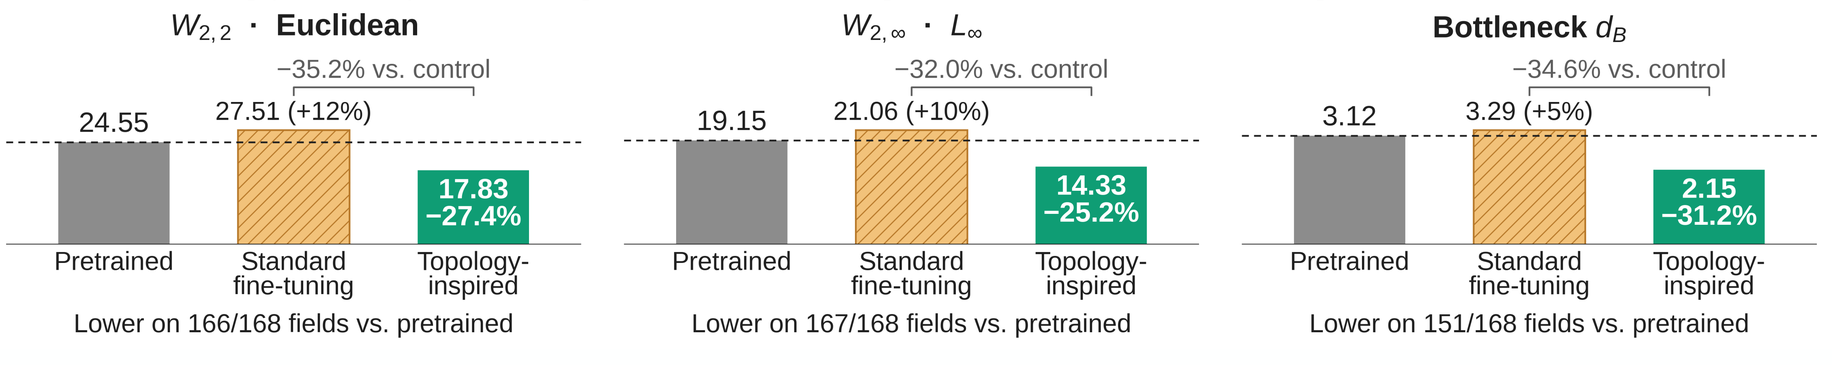

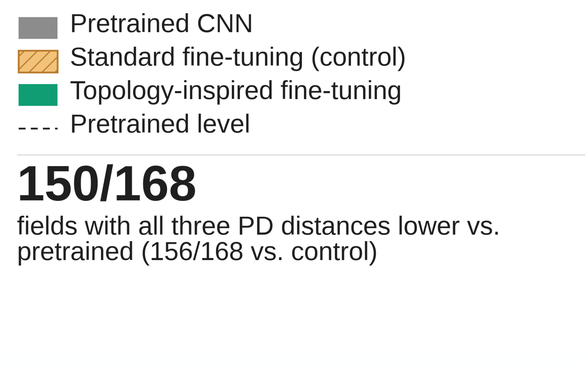

150/168
fields with all three PD distances lower vs. pretrained (156/168 vs. control)

Fidelity gains retained, broader structure improved

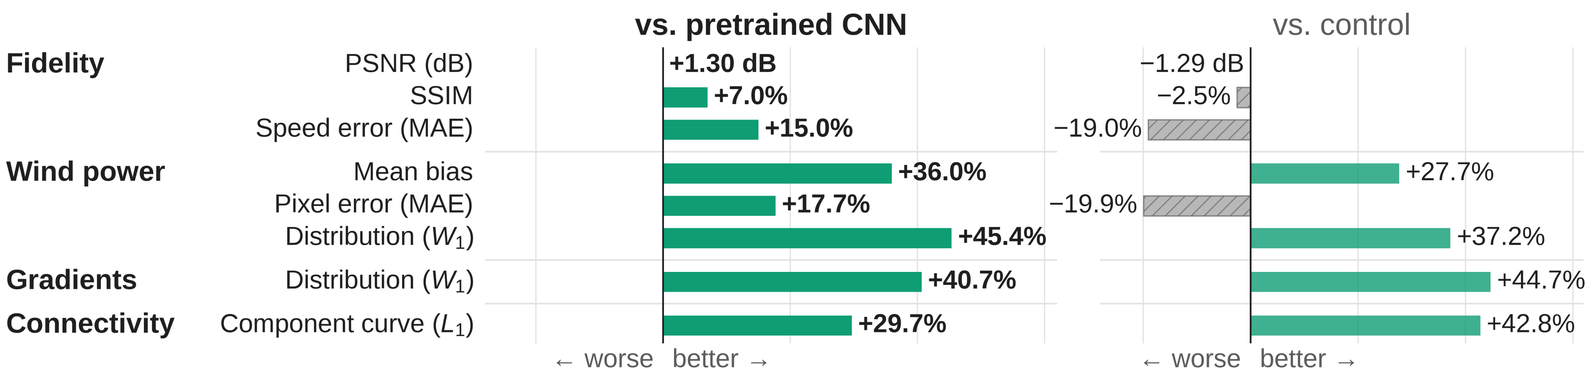

MT mostly agrees with PD

Of the fields where all three PD distances improve, how many also improve in MT?

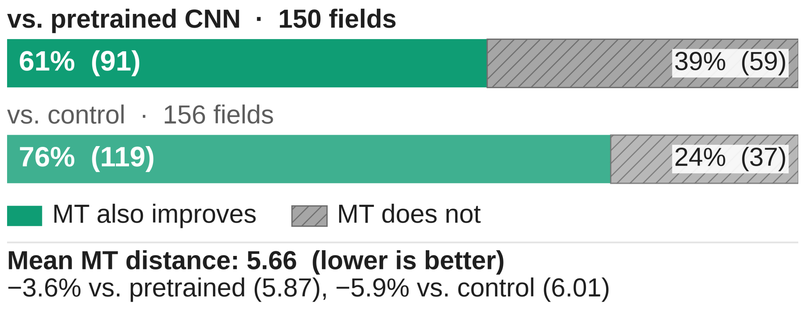

Mean MT distance: 5.66 (lower is better)
−3.6% vs. pretrained (5.87), −5.9% vs. control (6.01)

4. Qualitative Example

Sample 78: Conventional near-tie, better persistence agreement

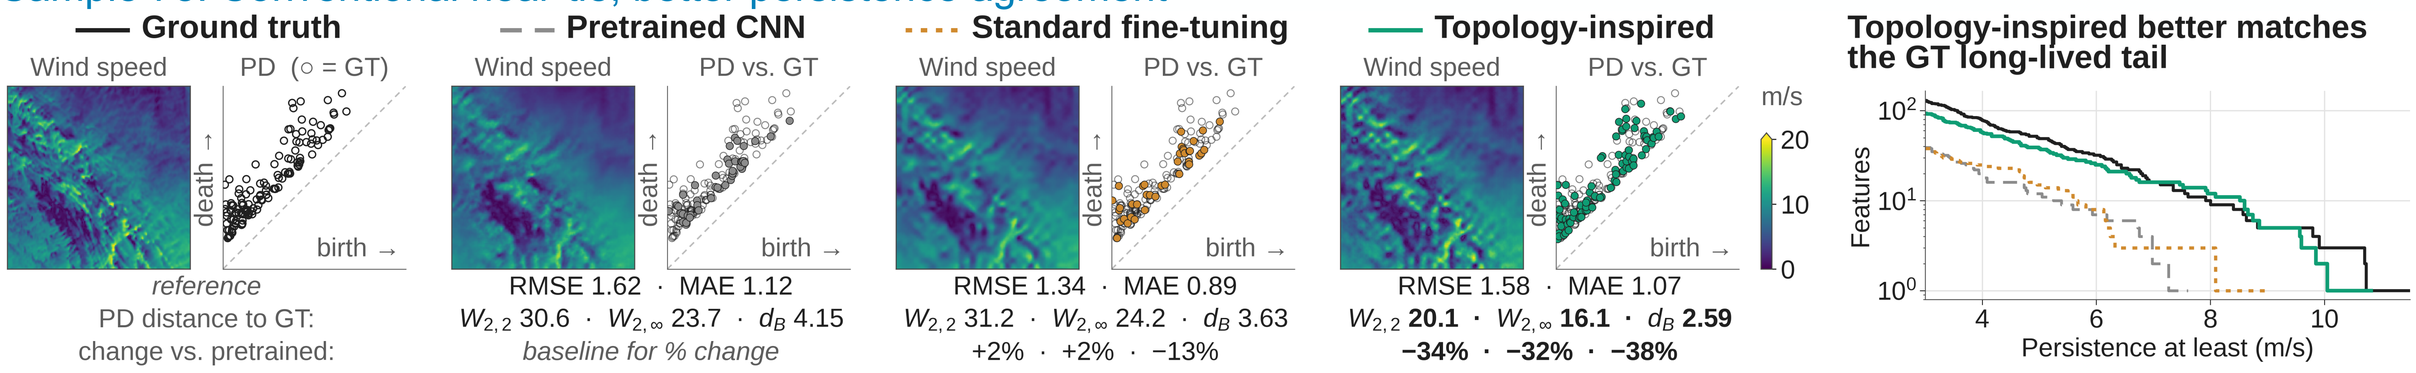

Standard fine-tuning has the lowest pointwise error; topology-inspired keeps pretrained-level fidelity yet lowers all three PD distances by 32–38% (plots: persistence ≥ 3 m/s; distances: all pairs).

5. Interpretation

Three results to remember

25–31%
lower PD distance vs. pretrained CNN

150/168
fields with all three PD distances lower vs. pretrained CNN

+1.30 dB
PSNR vs. pretrained CNN, fidelity retained

Limitations & Next Step

- PD/MT distances do not localize errors or directly measure physical displacement.
- The physical meaning and practical utility of the preserved persistent structures remain unclear.
- Gradient and persistence-pair terms were added jointly, so their separate effects are not isolated.

Next: map persistent features spatially; test wind-energy relevance across regions and SR models.

Full manuscript

arXiv:2610.22262
Methods · ablations
robustness · additional cases

Take Home: Topology-inspired fine-tuning substantially improves persistence agreement while retaining strong conventional fidelity; PD and MT reveal complementary structure.