


Learning from Neighboring Seismic Slices

Parameter-Efficient 2.5D Adaptation of DINOv3 for Seismic Denoising

Pablo Varela Bernal | pvarelabernal@tudelft.nl | TU Delft CSE3000 Research Project

01 Motivation

- Surveys capture subsurface reflections.
- Reflections form 3D volumes.
- Noise obscures weak geological structures.
- Structures continue across neighbouring slices.
- 2D models ignore cross-slice context.
- 2.5D uses nearby slices as extra input.
- Goal: better denoising with low adaptation cost.



02 Research Question

Does 2.5D spatial context improve parameter-efficient adaptation of pretrained visual foundation models for seismic denoising compared with standard 2D adaptation?

04 Dataset

Training

- ThinkOnwards Impeccable
- Synthetic
- Time slices

Robustness

- Netherlands F3
- Real Survey
- Inline/Crossline slices

05 Metrics

MS-SSIM

- Structural
- Higher is Better

MS-SSIM-R

- Residual
- Lower is Better

MSE


- Pixel Error
- Lower is Better

08 Conclusion


- 2.5D context improves synthetic denoising:
 - 3-channel 2.5D gives the best accuracy-cost trade-off.
 - 5-channel 2.5D gives the strongest synthetic reconstruction.
- Field transfer on F3 depends on orientation:
 - Time neighbours are weakly correlated, so 2.5D uses misleading context and over-smooth the output.
 - Inline/Crossline are strongly correlated, so 2.5D provides useful context and preserves more structure.

03 2D vs 2.5D DINOv3 Denoising Pipeline

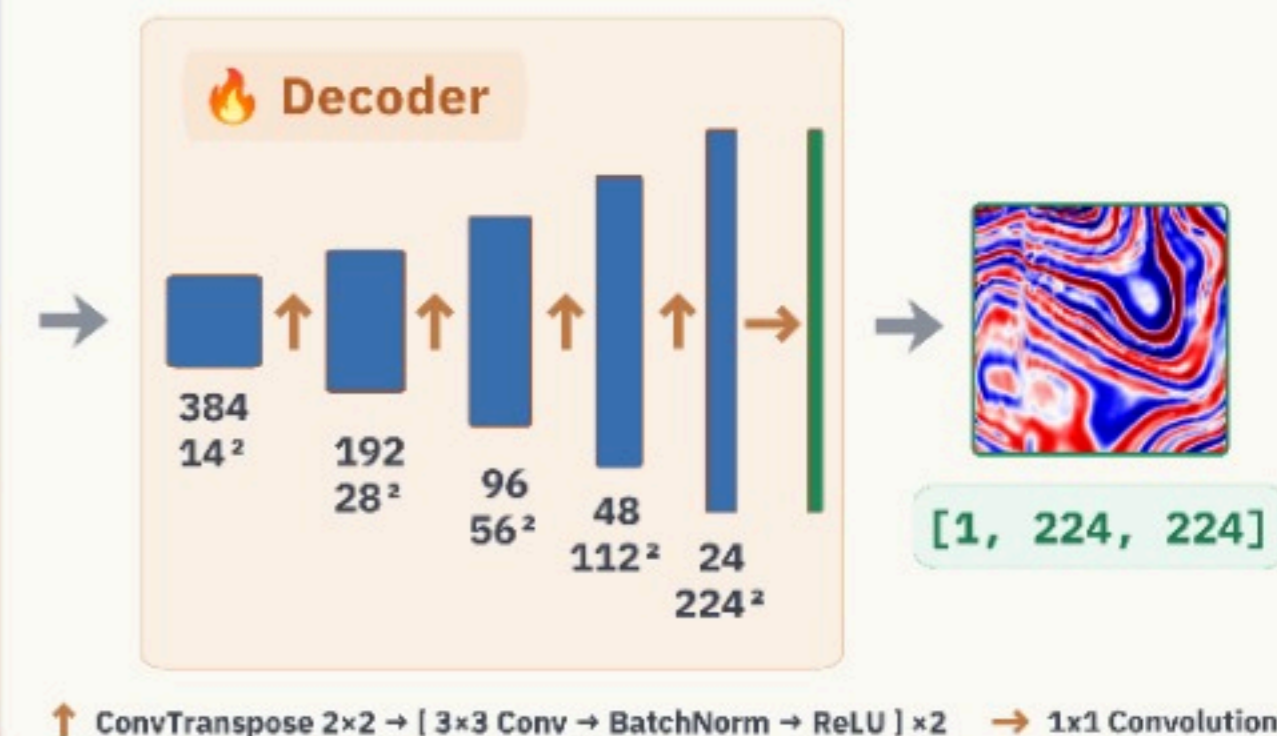
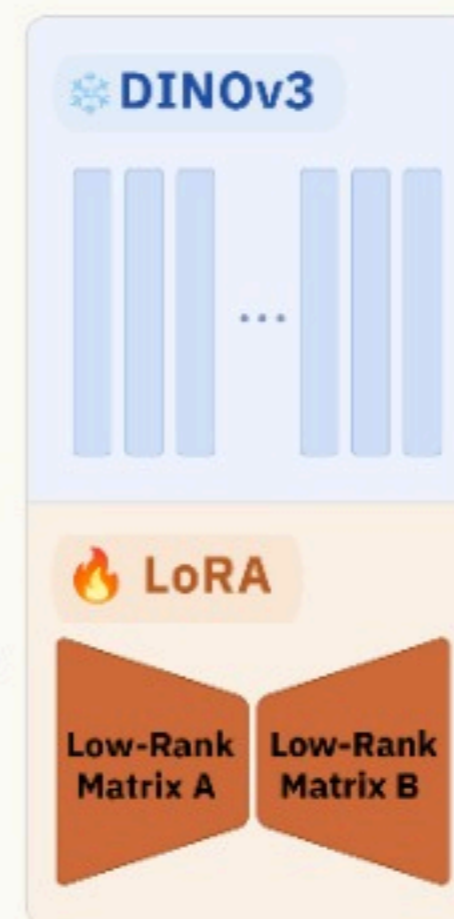

2D · 1-channel



2.5D · 3-channel



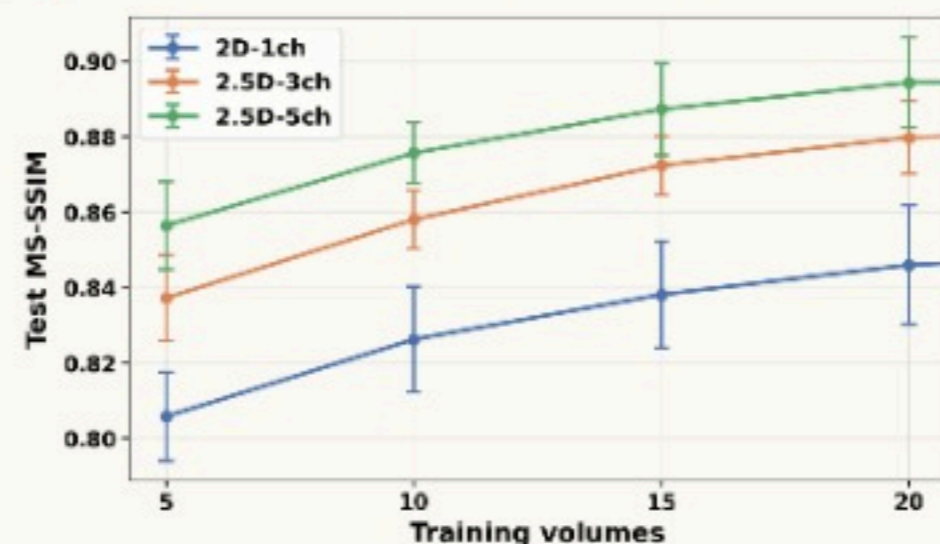
2.5D · 5-channel



🔥 Trainable

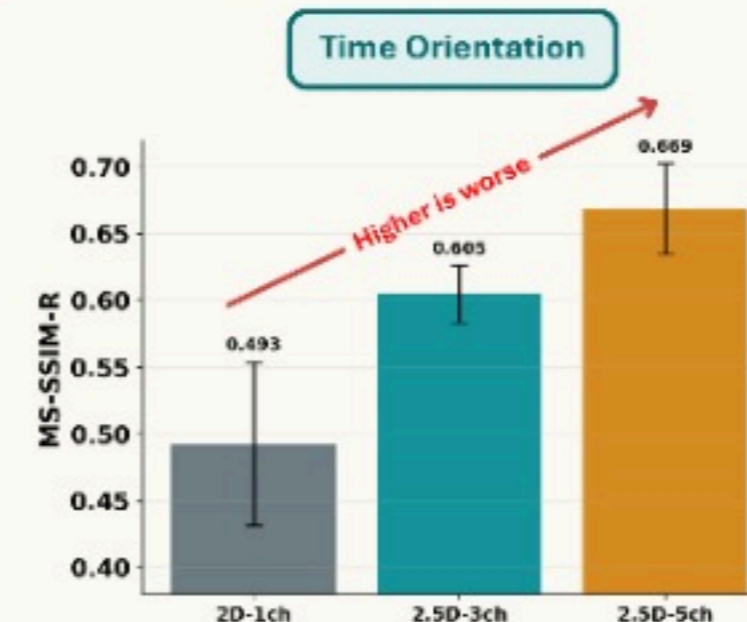
❄️ Frozen

06 Denoising on Image Impeccable

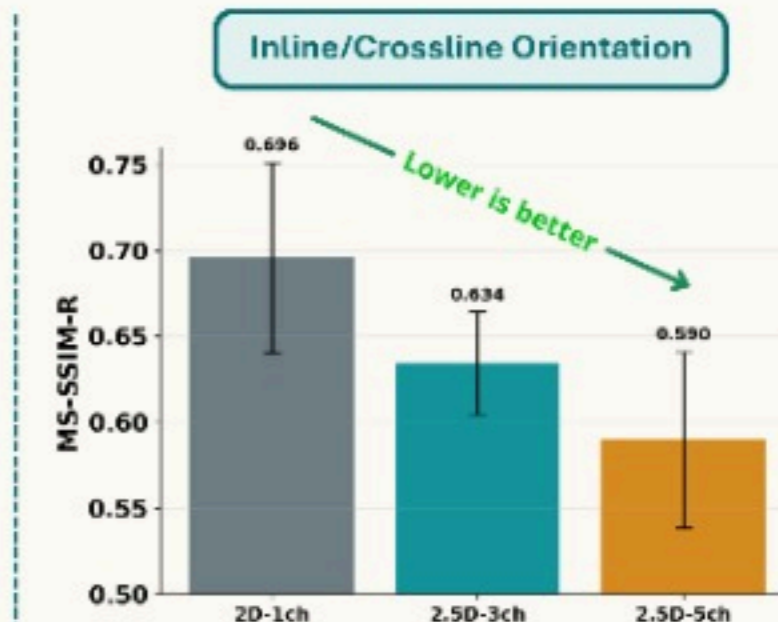


2.5D achieves more accurate denoising with less training data.

07 Denoising on F3



2.5D oversmooths the output and removes structure



2.5D preserves finer detail and removes less structure

