


Where Does Adaptation Matter?

Layer-wise Importance of Parameter-Efficient Adaptation of Vision Foundation Models to Seismic Denoising

Okan Demir Baykal | Supervisors: Dr. Jing Sun, Dr. Eric Verschuur, Dr. Tiexing Wang, Jiahua Zhao | CSE3000 Research Project

Problem

- **Seismic denoising** is the process of removing acquisition noise from seismic recordings to recover coherent reflection events.



- Recent work has shown that **parameter-efficient fine tuning (PEFT)** of powerful vision foundation models, pre-trained on natural images, is a viable method for denoising.


Question

- PEFT is often applied uniformly across model layers, but a ViT's layers are not uniform in representation: early layers handle low-level structure, later ones abstract semantics.

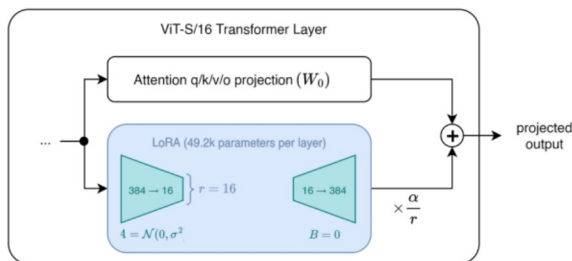
Research Question: Which layers are most important to adapt?

- To study this, we place two structurally different types of PEFT methods to different layers of a ViT backbone.

Pfeiffer Bottleneck Adapters




Low Rank Adaptation



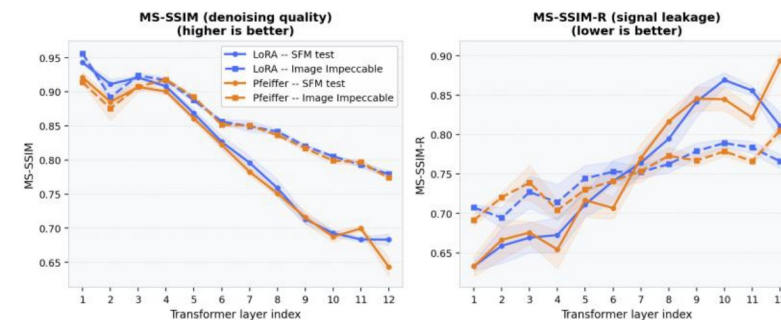
Layer-wise performance profile

- Trained and tested on an active-source dataset from Sheng et al.
- Primary metric is Multi-Scale Structural Similarity (MS-SSIM) between the original image, the denoised image, and the residual.



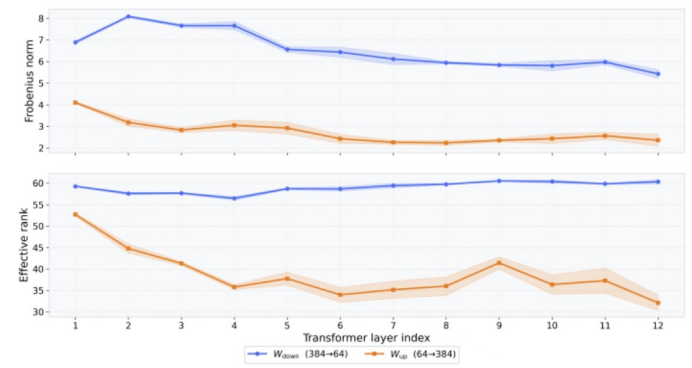
A clear **early-to-late importance ordering**: Even adapting just the first layer is competitive with the performance of full adaptation

Generalization Probe



Comparing results from the Sheng et al dataset and the Image Impeccable denoising challenge dataset shows agreement in the layer-wise performance profile.

Pfeiffer Weight Analysis



Analyzing the weights of the fully adapted Pfeiffer model shows a similar, but less clear pattern.