


# Improving Generalization Of VFMs On Denoising Seismic Data Through The GRQO RL-based Fine-Tuning Method

Andrei Oprescu  
TU Delft BSc Computer Science and Engineering




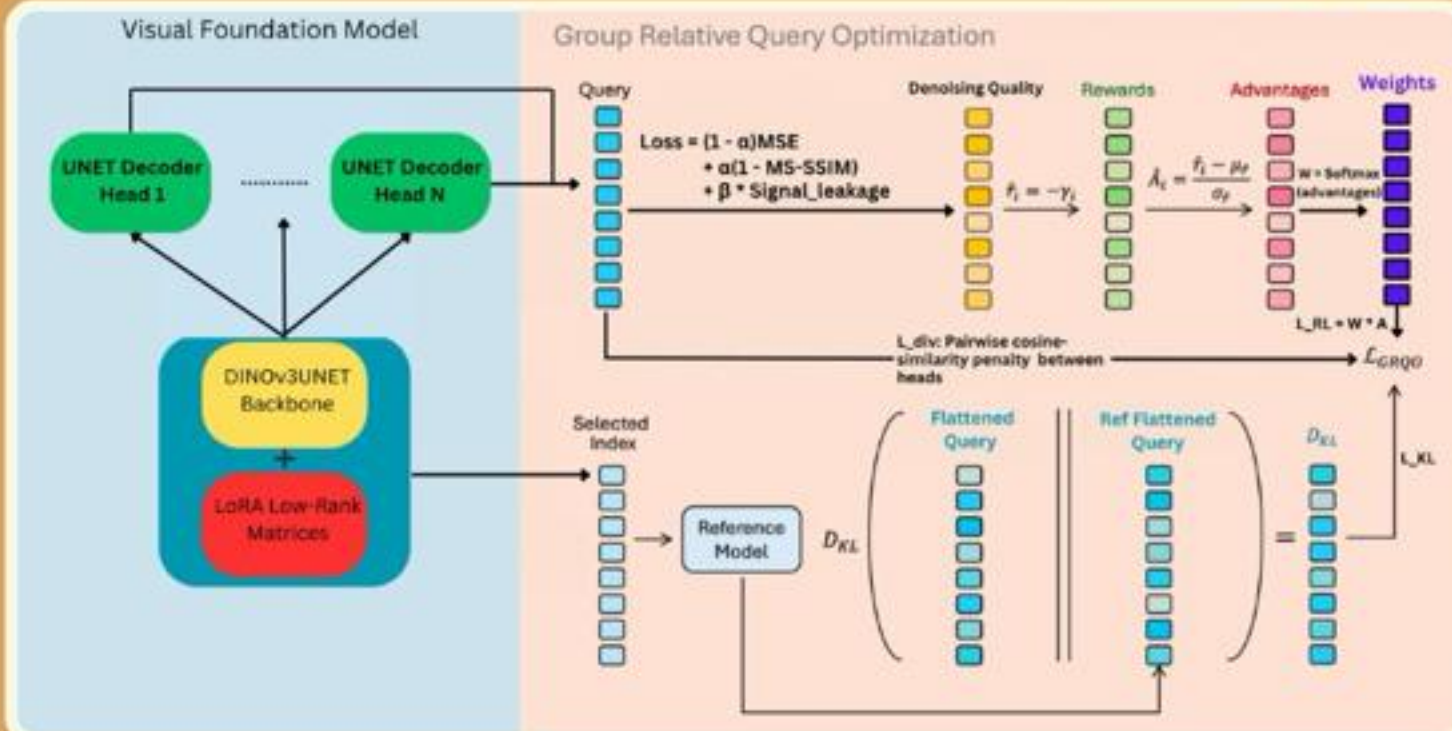


## Task: Efficient Seismic Denoising

- Remove noise
- Retain signal
- Recent solutions use VFMs with TTA to generalize to new domains at inference time



**Problem:** even though Test Time Adaptation (TTA) works, it's slow!  
**Goal:** generalise across domains with GRQO, and get rid of (or reduce) TTA training.





### Datasets

- SFM synthetic data
- Image Impeccable synthetic data
- SFM real data

### Metrics

- MS-SSIM
- MS-SSIM-R
- LS
- PSNR



### Conclusion

#### Generalization of GRQO - UMAP Analysis

- GRQO improves signal retention at near-zero inference cost
- GRQO reduces but cannot replace TTA.

Figure: DINOv3 encoder embeddings (mean-pooled), UMAP-projected, coloured by domain and kurtosis class. Left: SFT-20. Right: SFT-20 + GRQO-10.

- Synthetic train + test → one tight island
- Real SFM → isolated cluster, low-kurtosis only.
- Left ≈ Right: GRQO leaves the geometry unchanged.