

# Seeing Through Seismic Noise with Soft Spatial Blending

## Parameter-Efficient Spatial Blending of Vision Foundation Models for Seismic Denoising

Student: Alexis FIMEYER, A.H.P.A.FIMEYER@student.tudelft.nl


Supervisors: Dr. Jing Sun, Dr. Tiexing Wang, Dr. Eric Verschuur, Jiahua Zhao

Course: CSE3000



### 1 Problem: Seismic Denoising

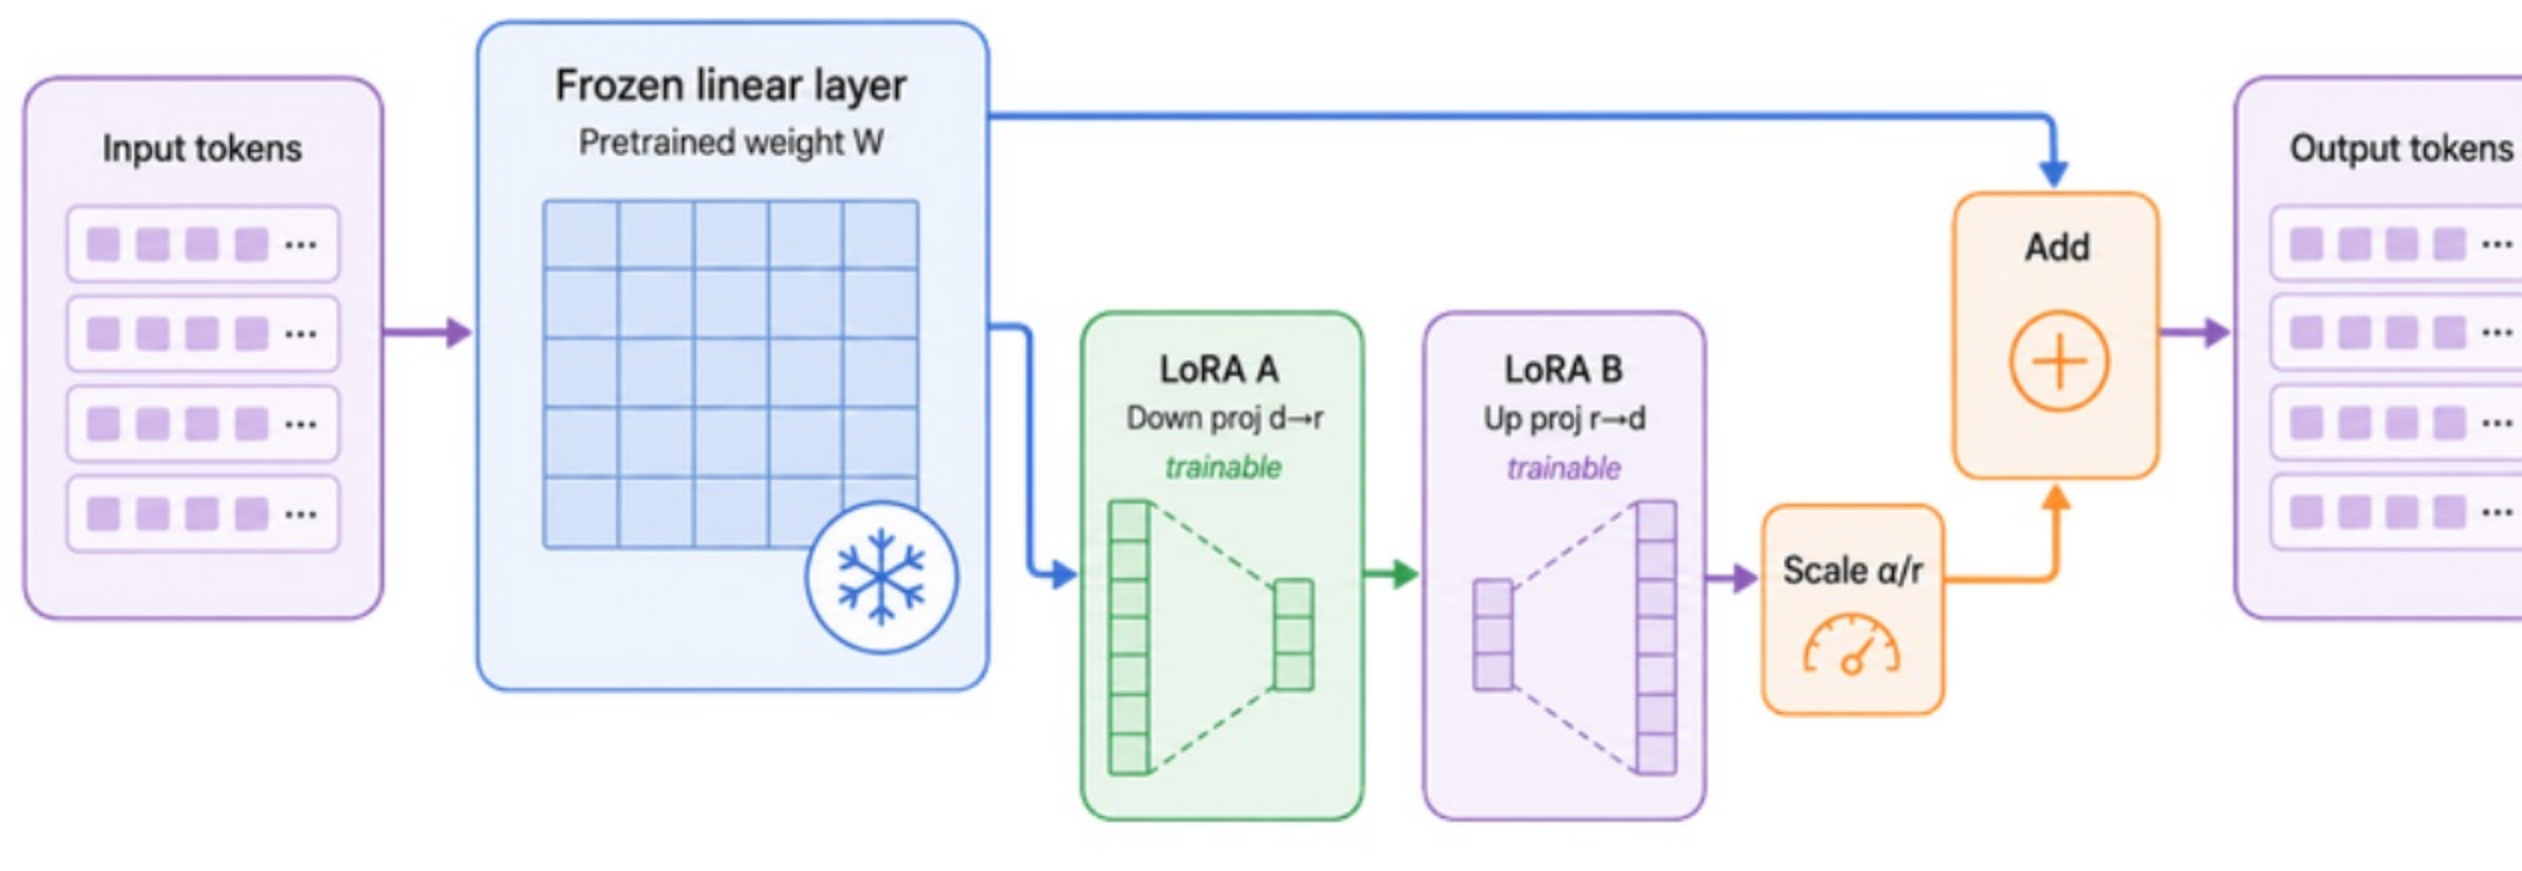
- Remove noise
- Preserve reflector continuity
- Avoid smoothing weak structures



Denoising is only useful if geological structures remain visible

### 2 LoRA Experts

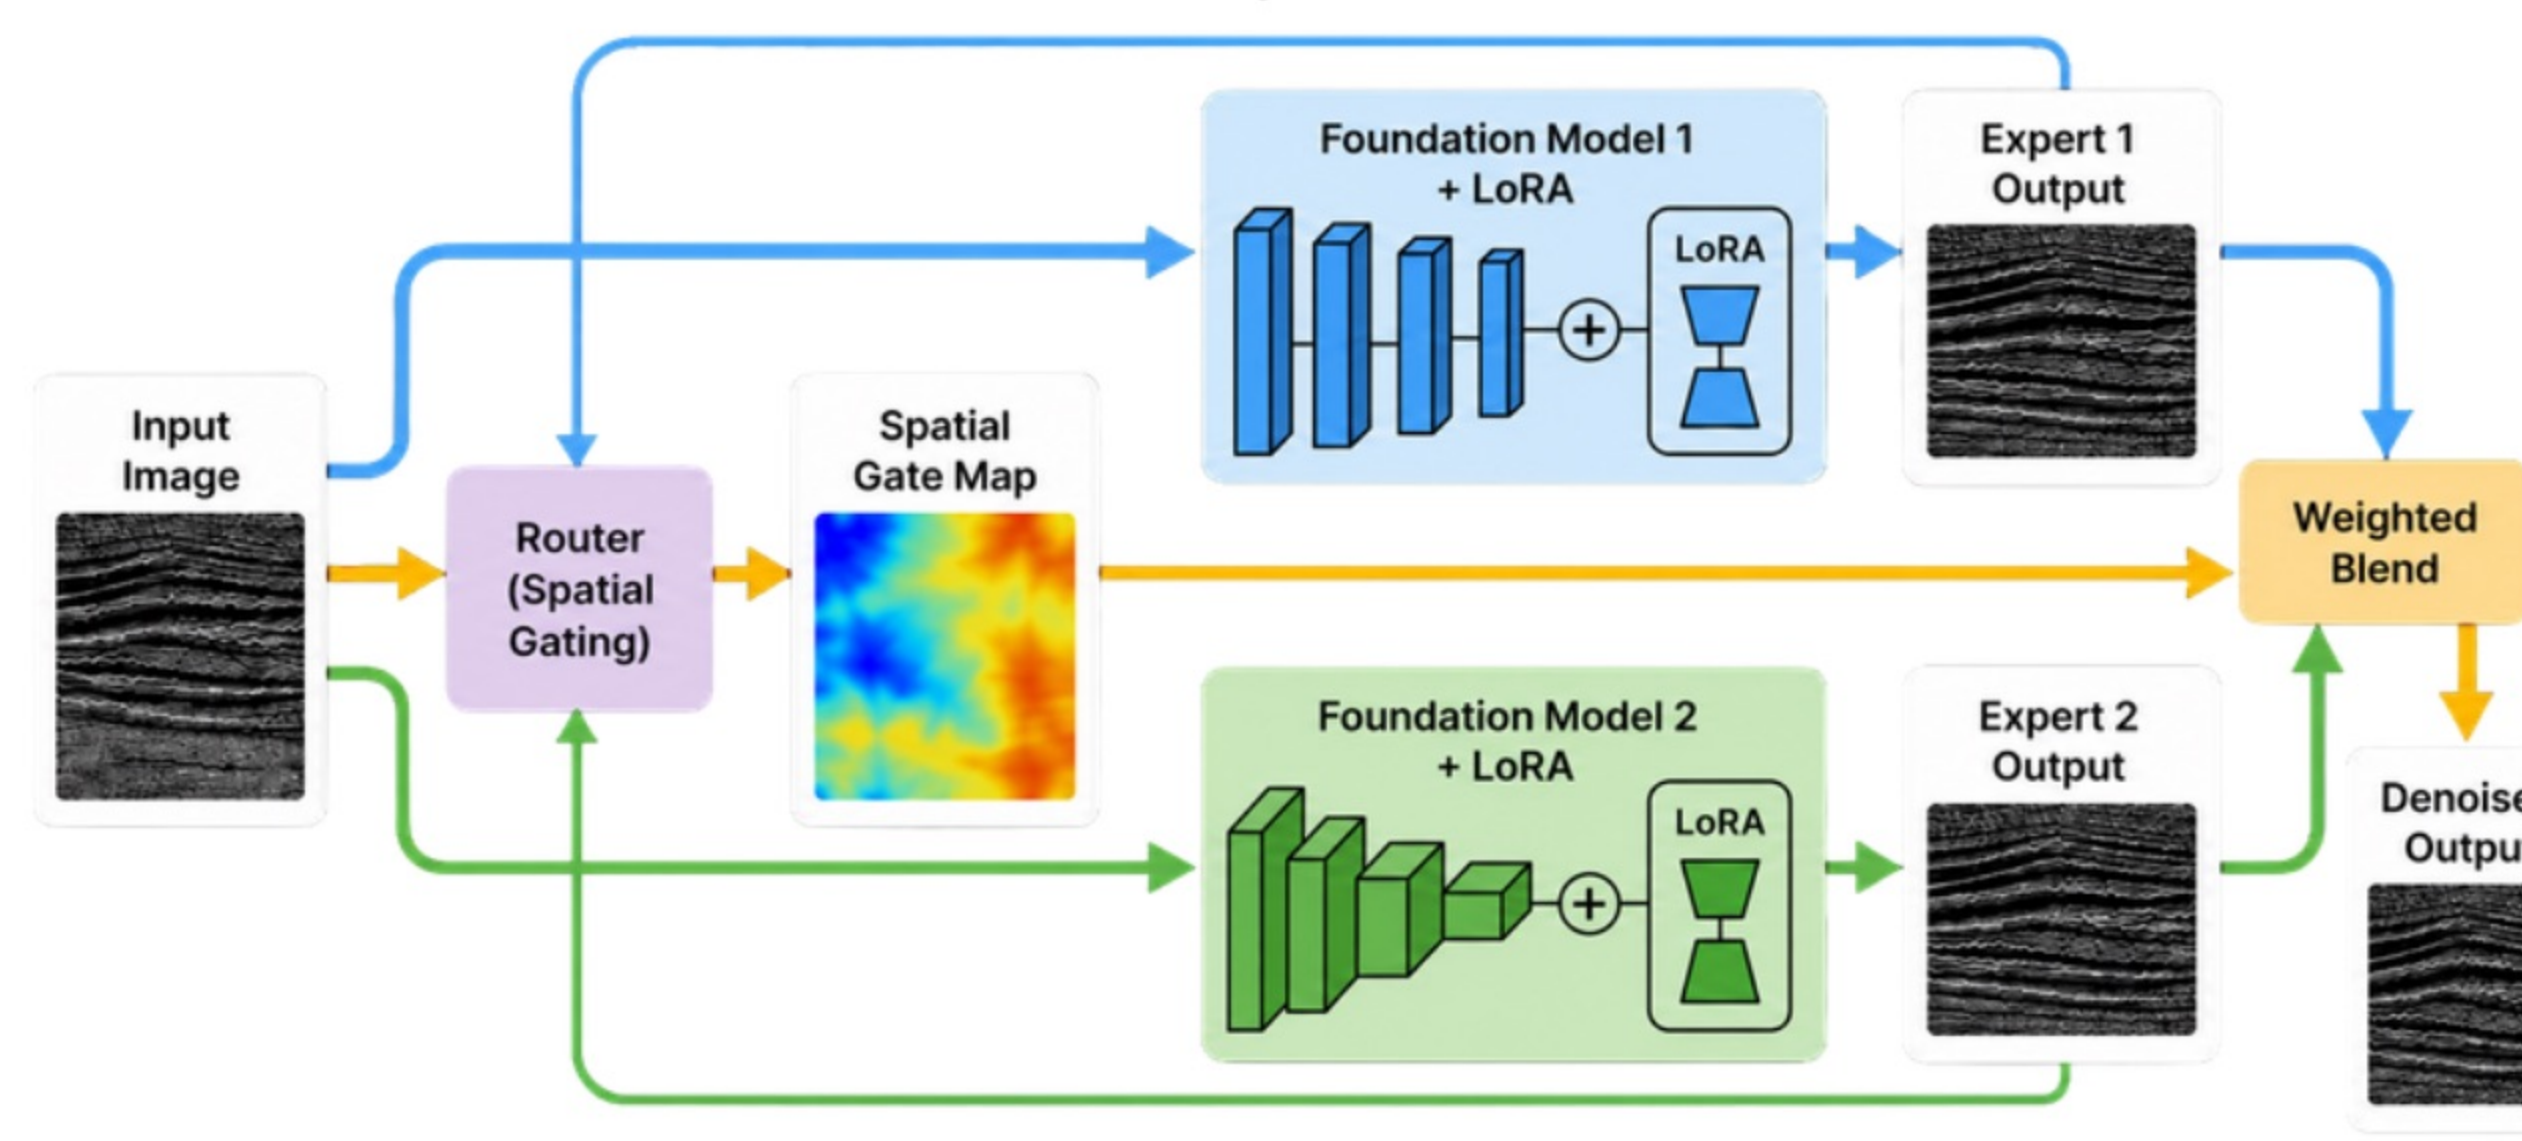
- Trained on 1600 pairs
- 100 Epochs
- Rank  $r = 16$ ,  $\alpha = 64$
- Improvement from base on SSIM:
  - DINOv3: 172.9%
  - Swin v2: 159.7%



DINOv3 is the stronger standalone LoRA expert model

### 3 Soft Spatial Blender Pipeline Architecture

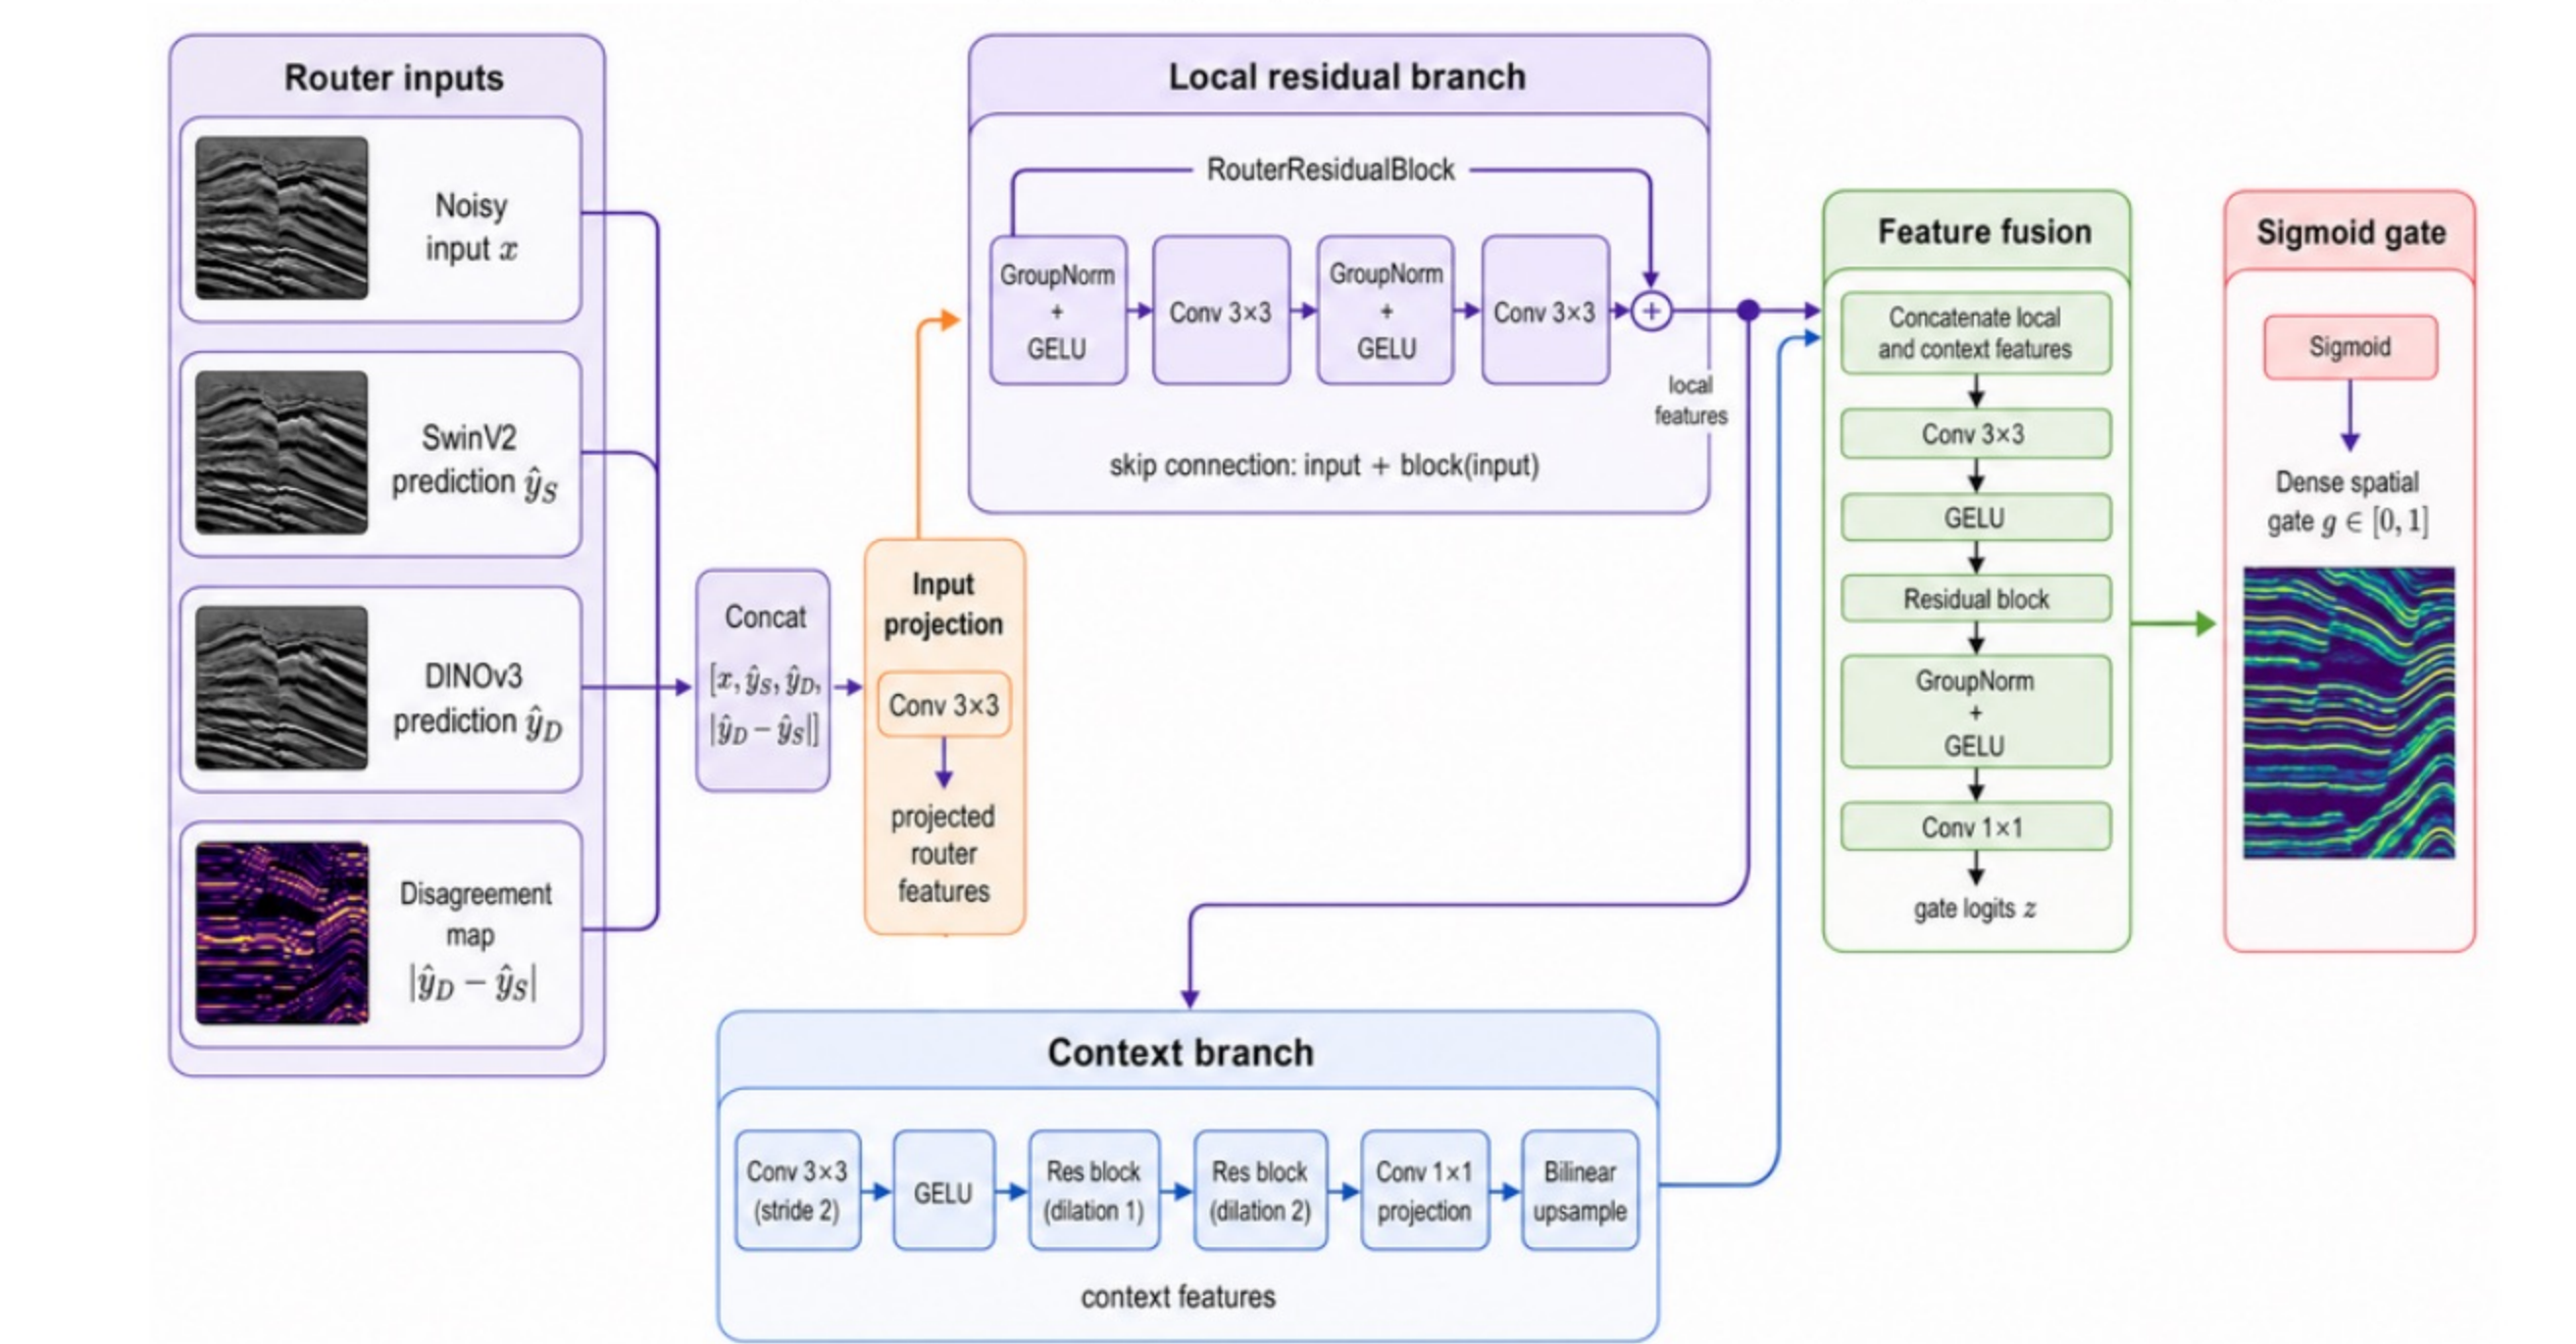
- Output =  $(1 - g) \text{Expert 1} + g \text{Expert 2}$
- Router predicts per-pixel gate
- Router sees where experts disagree
- Blends instead of choosing one model



The router turns two denoisers into a spatial mixture

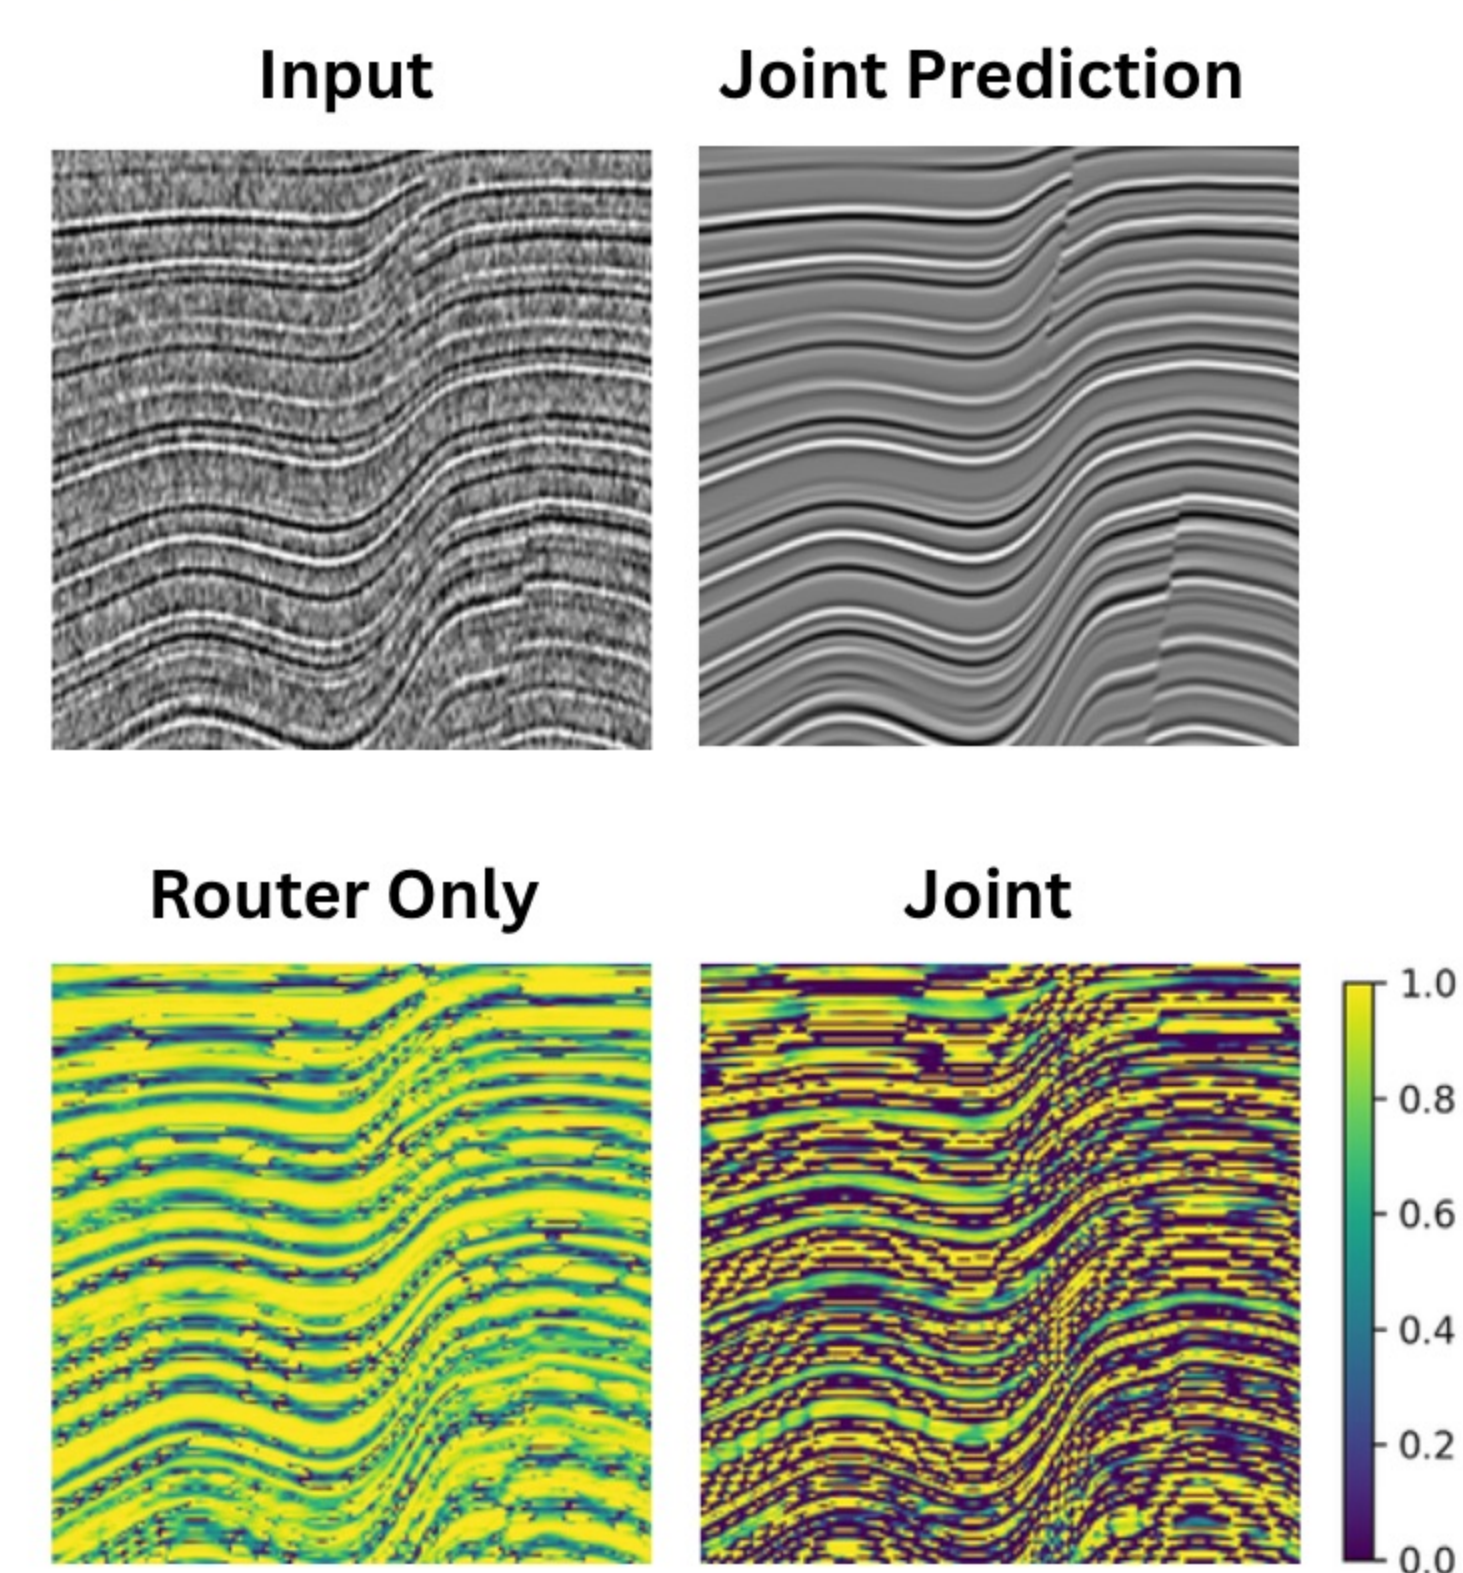
### 4 Router Architecture (How to Predict $g$ )

- Loss function:  $L = L1 + 0.005 \text{TV}(g)$
- $\text{TV}(g) = \text{mean } |g[i+1, j] - g[i, j]| + \text{mean } |g[i, j+1] - g[i, j]|$



The router creates the gate map using the generated  $g$  values

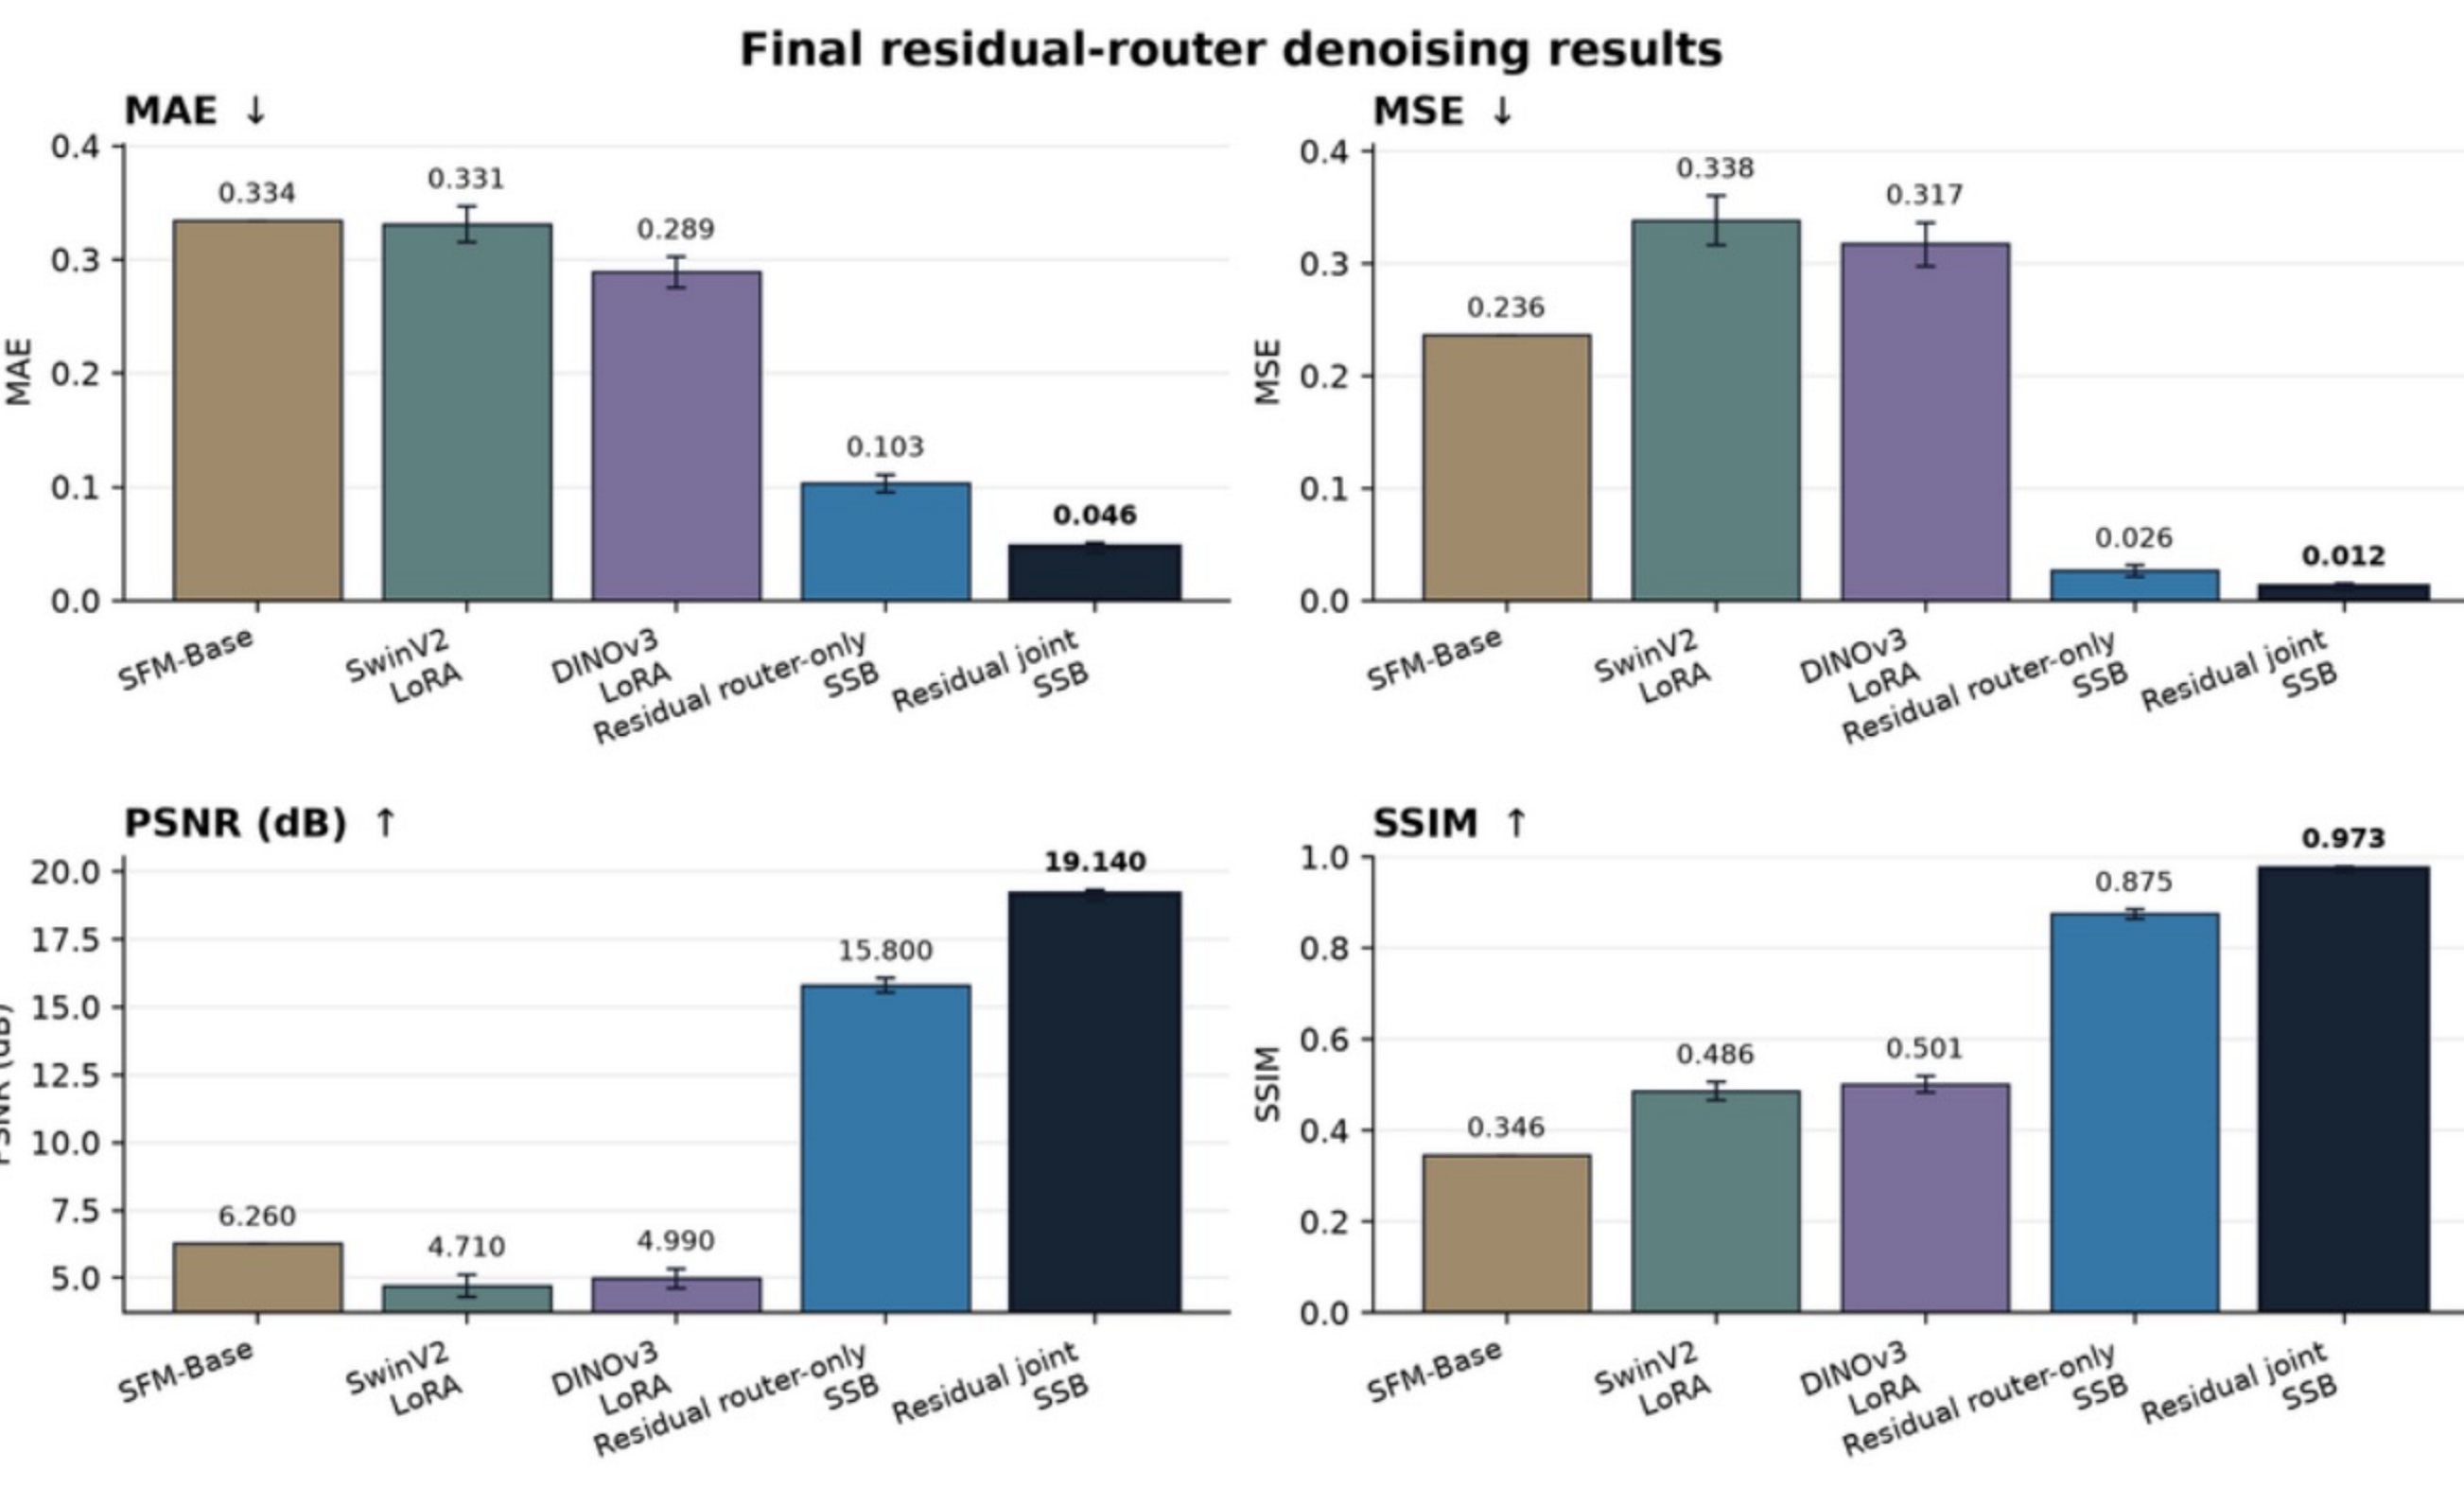
### 5 Joint Training of Experts and Router



- $g = 1 \rightarrow$  DINOv3
- $g = 0 \rightarrow$  Swin v2
- Router Only prefers DINOv3 71.3% of the time
- Joint choses it only 32.7% of the time
- Expert aware training switches the role of experts

Joint training inverses the role of the experts

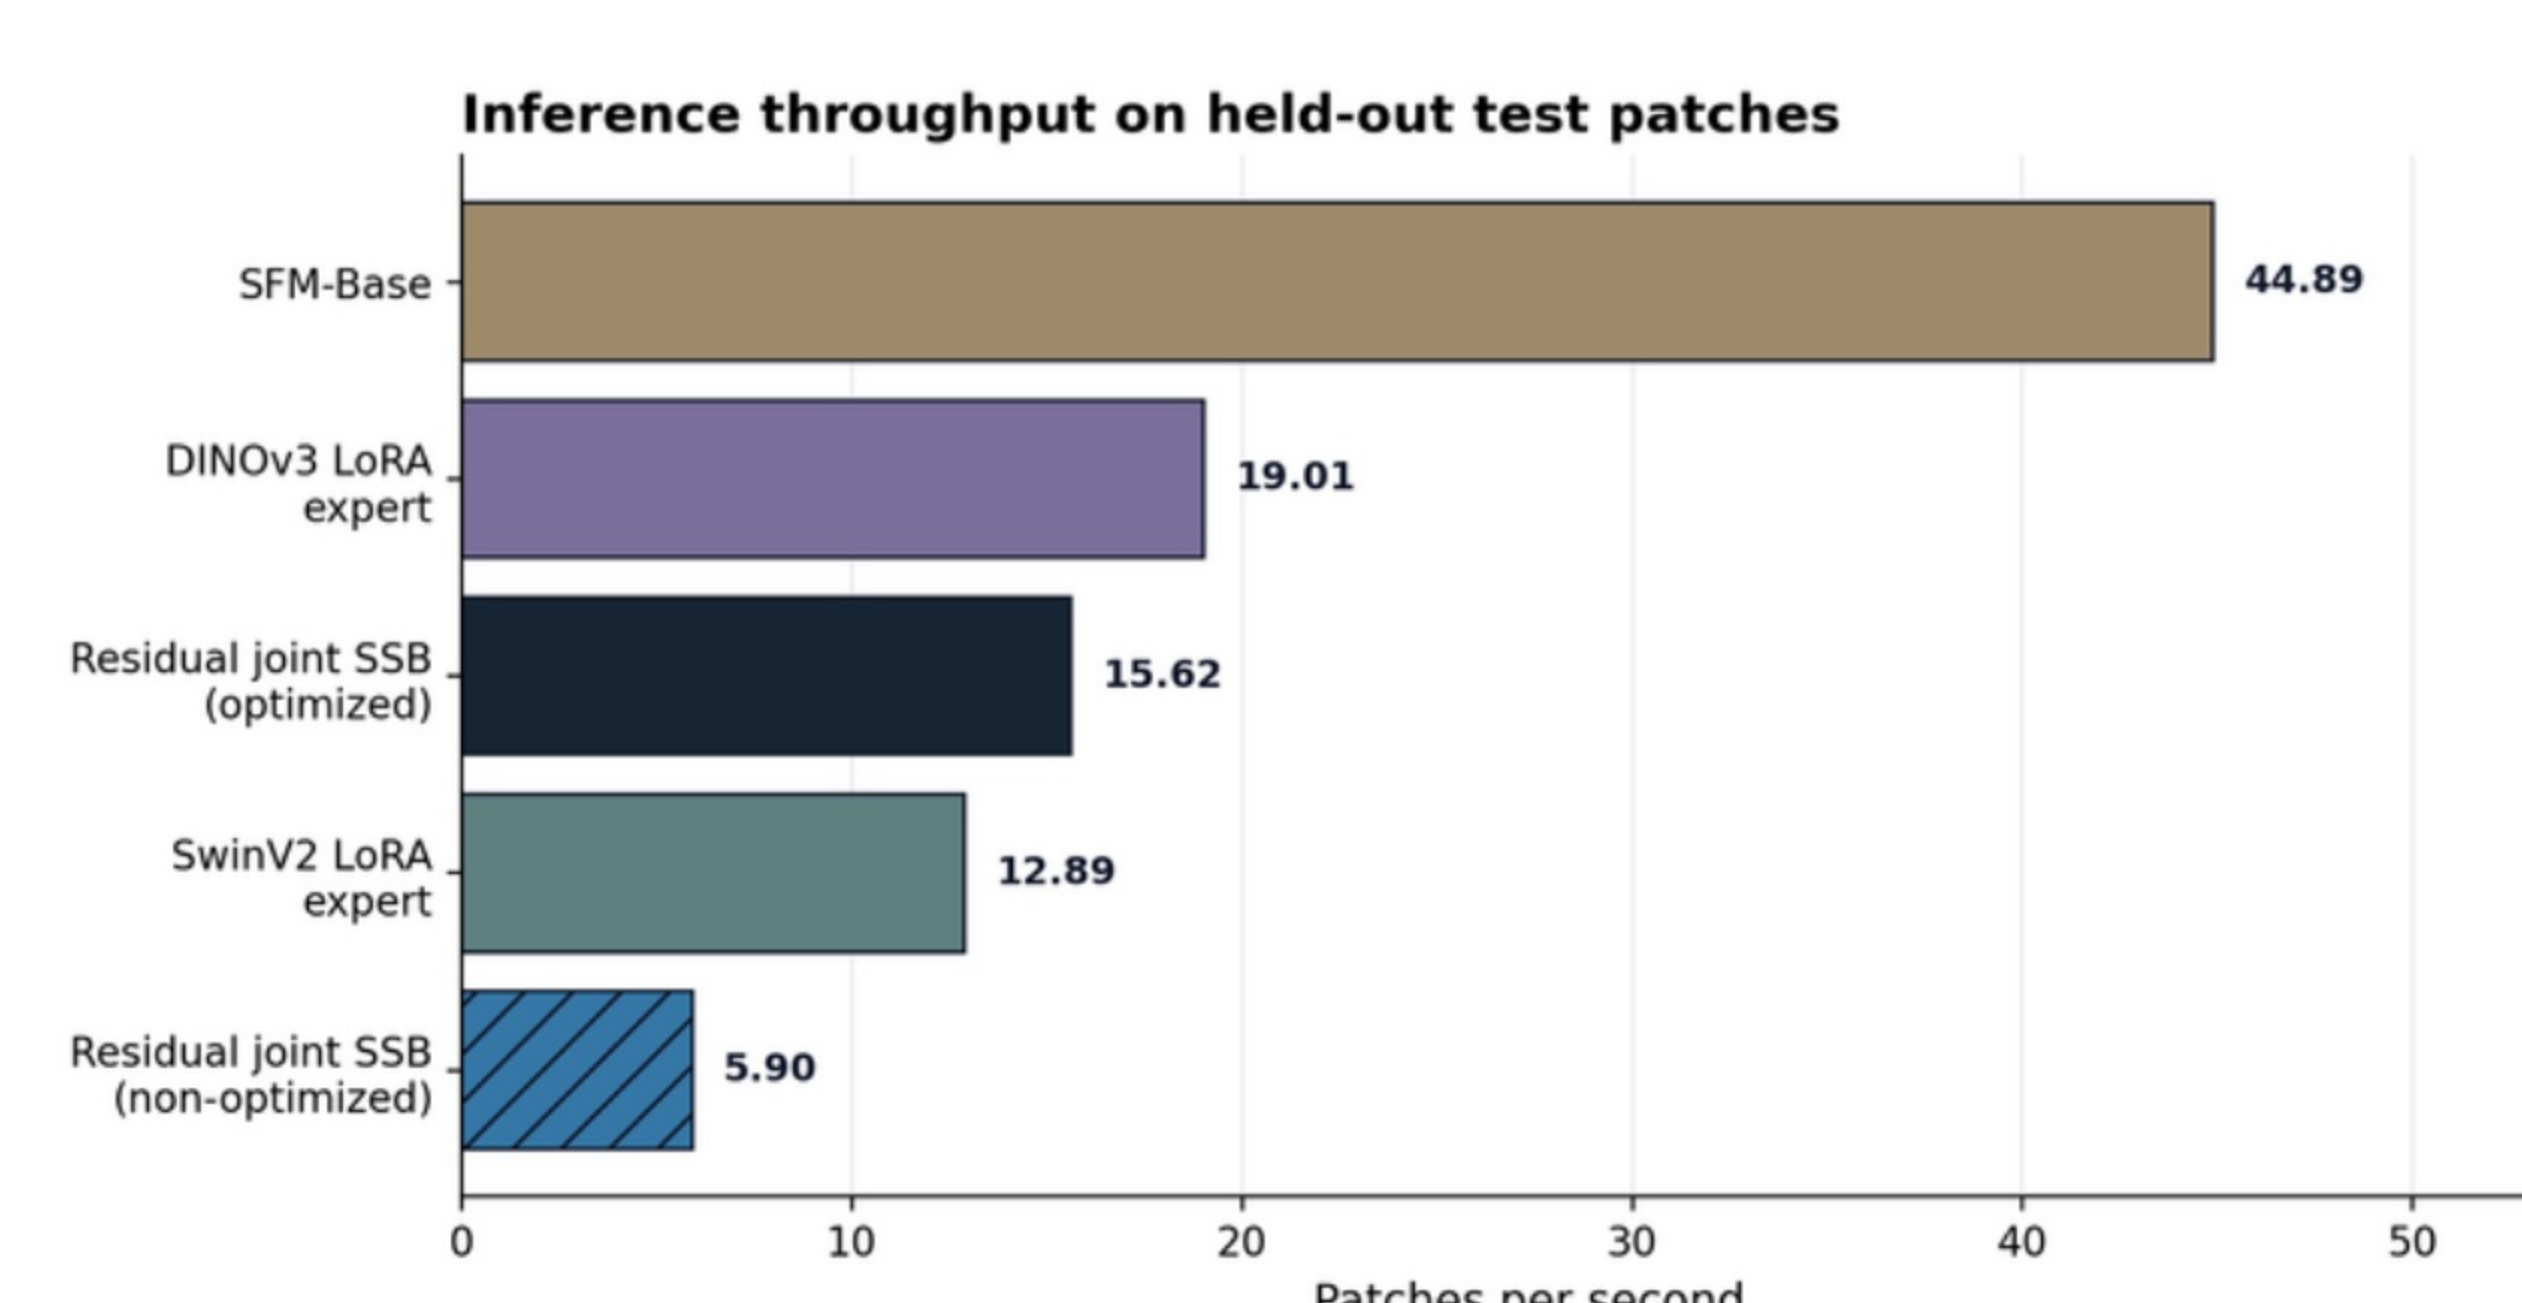
### 6 Final Results and Evaluations



Best model: jointly trained residual SSB

### 7 Inference Speed

SSB Evaluates two experts for each prediction  
Trade-off between inference speed and denoising quality



Trade-off to be made between: Inference Speed vs Quality

### 8 Main Contributions

- 1 LoRA adaptation of SwinV2 and DINOv3
- 2 Adaptable SSB pipeline
- 3 Router Architecture
- 4 Trained Model
  - Beats SFM Base
  - Can be adapted efficiently

Jointly SSB improves denoising quality, at extra inference cost