Poster #567. Poster session #1.
ExHall A-F, June 5, 2026.

Project Page

SNP INRs: The Surprising Effectiveness of Noise Pretraining for Implicit Neural Representations

Kushal Vyas[1], Alper Kayabasi[2], Daniel Kim[1],
Vishwanath Saragadam[2], Ashok Veeraraghavan[1], Guha Balakrishnan[1]

[1]Rice University

[2]University of California, Riverside

CVPR JUNE 3-7, 2026 DENVER COLORADO

Implicit neural representations (INRs)

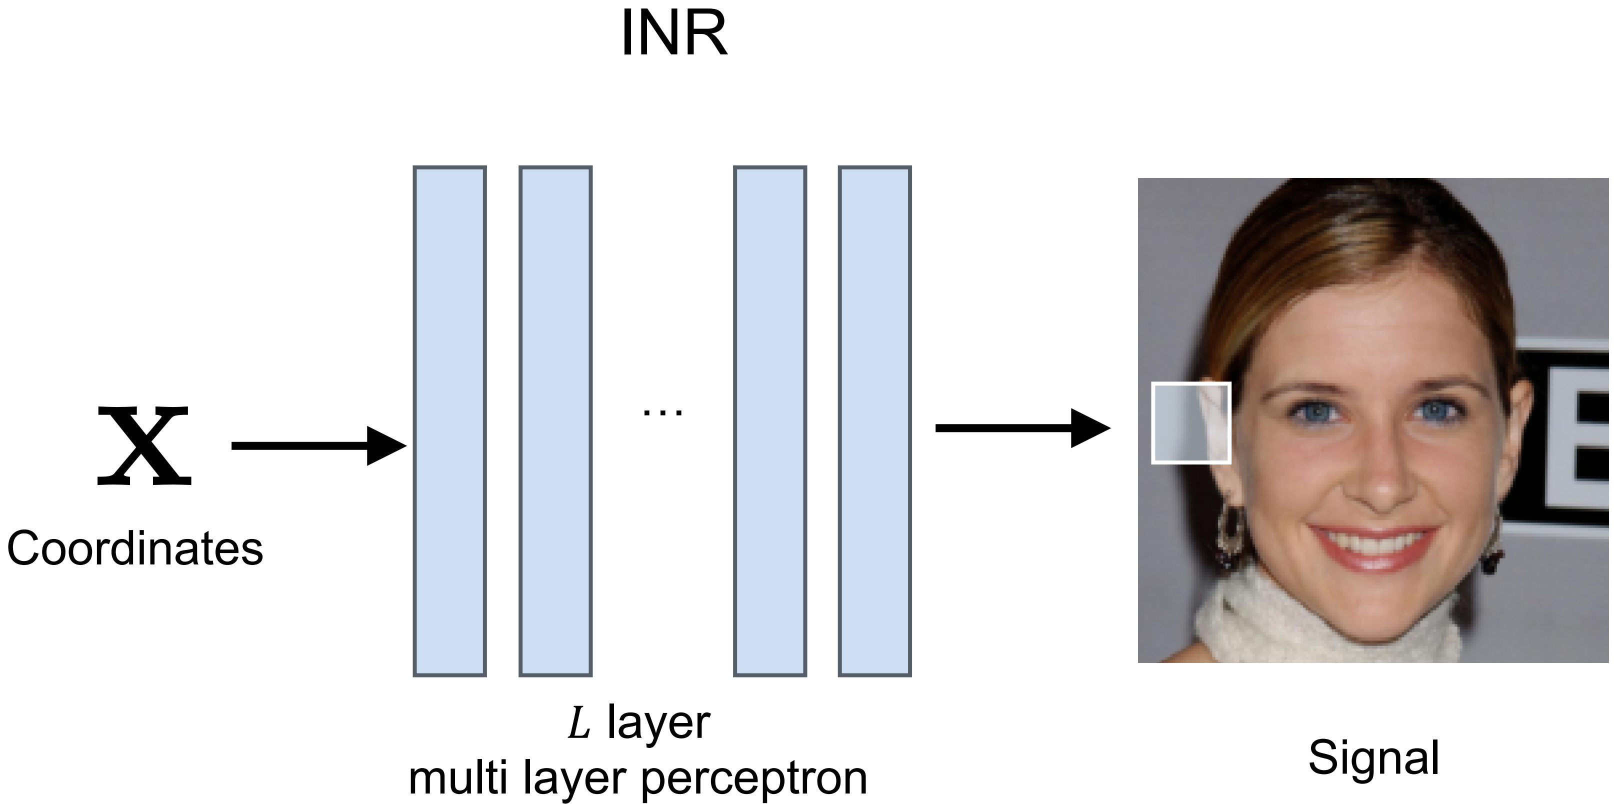

Random init. INRs lack generalizable features

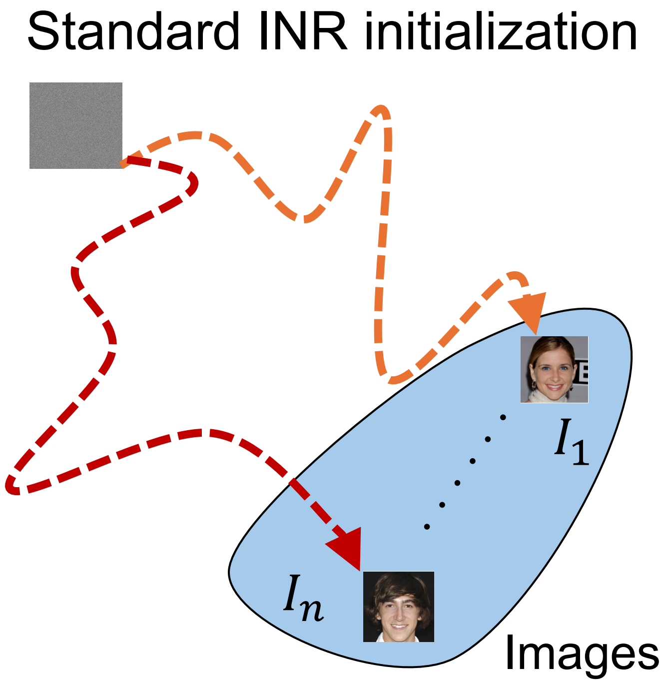

INR optimization trajectory
converging to target images

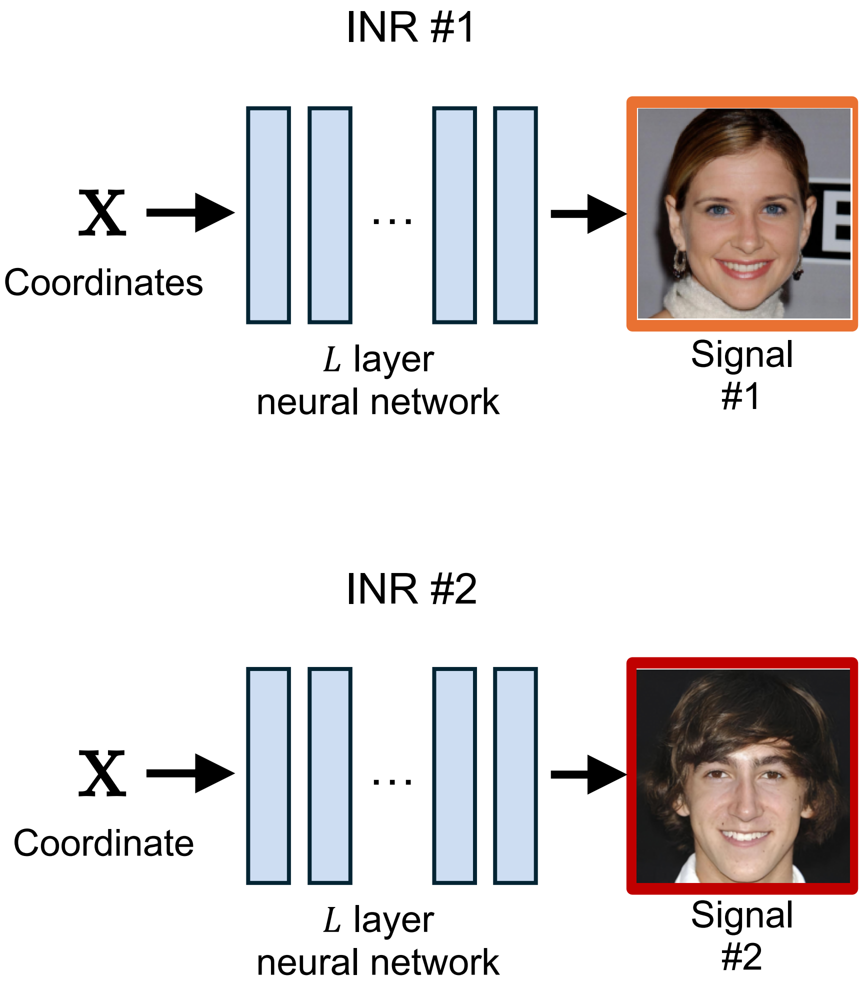

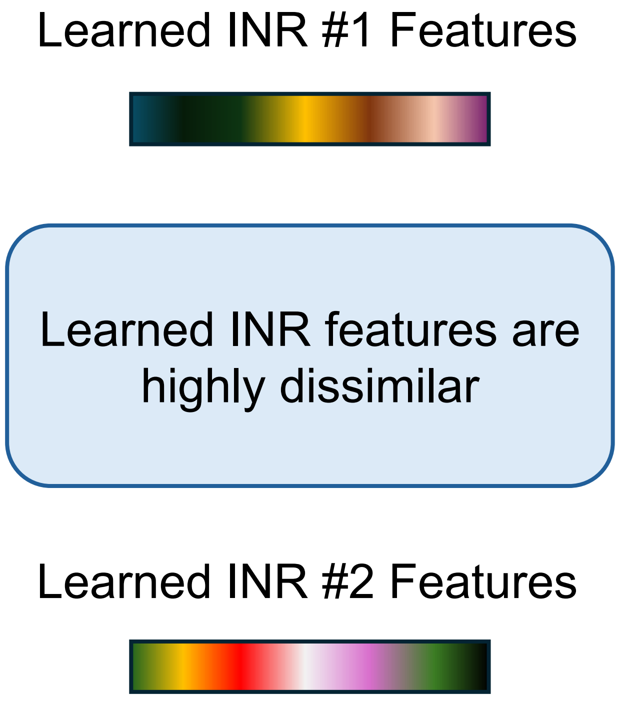

Learning data-driven INR initializations

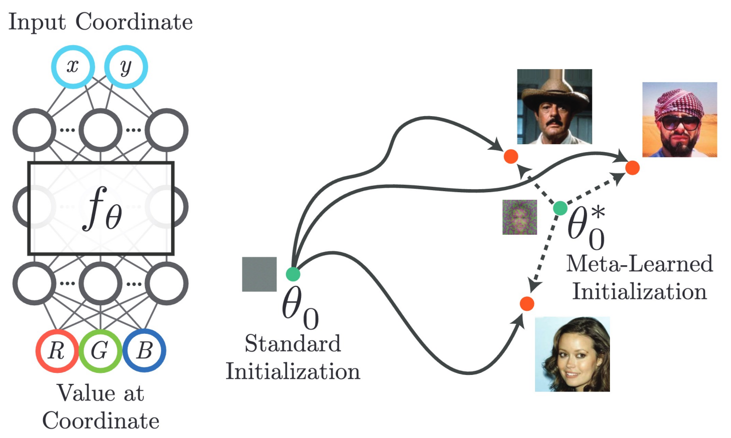

Tancik, Mildenhall et.al., LearnIt, CVPR 2021

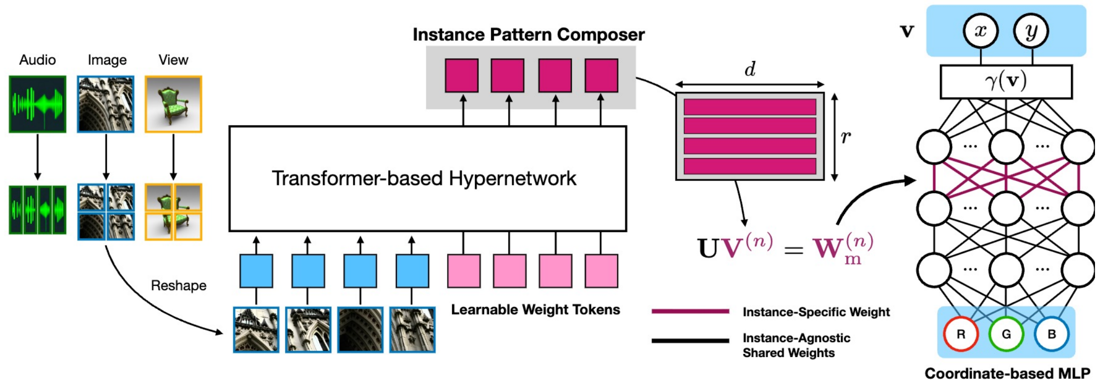

Kim, Lee et.al., Implicit Pattern Composers (IPC), CVPR 2023

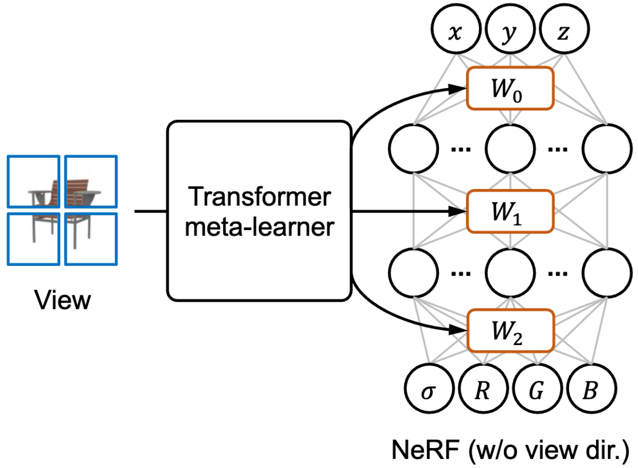

Chen et.al., TransINR, ECCV 2022

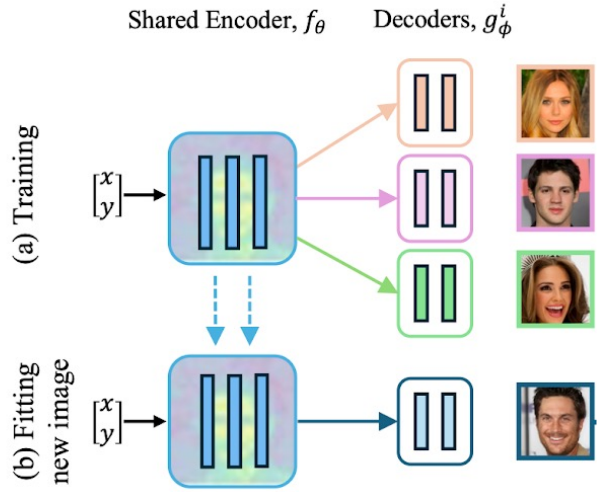

Vyas et.al., STRAINER, NeurIPS'24

Prior work: Transferable INR features

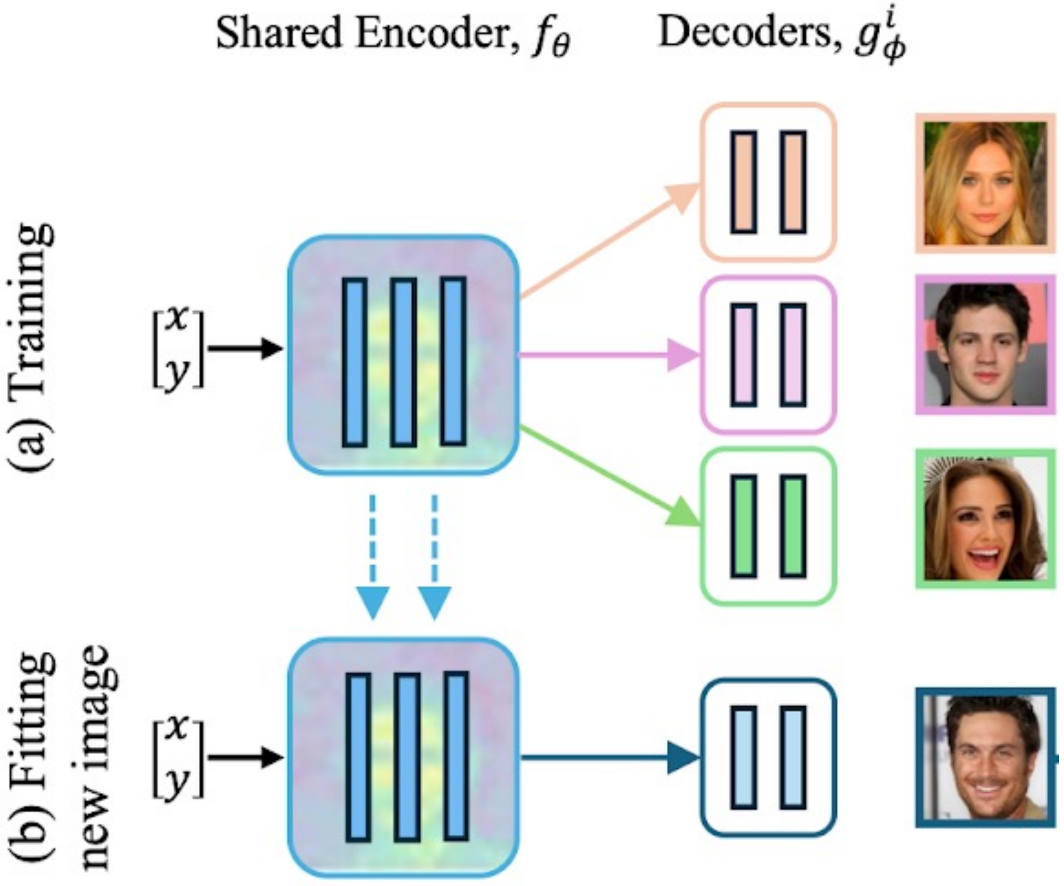

Vyas et.al., STRAINER, NeurIPS'24

- **Pretraining (learning INR initialization):** STRAINER shares initial layers in the INR to capture **low-frequency structural details** while jointly fitting multiple signals of the same class.

- STRAINER **pretrained encoder** serves as a powerful and generalizable initialization for INRs

- **Test-time fitting on unseen signal:** STRAINER **pretrained encoder is retained,** along with a randomly initialized network. An unseen test signal is then iteratively fit to the INR.

Does the benefit of data-driven initialization stem from dataset-specific mid-to-high-level features, or simpler universal natural image statistics?

Does the benefit of data-driven initialization stem from dataset-specific mid-to-high-level features, or simpler universal natural image statistics?

Can we learn high-quality INR initializations from just image statistics, without any actual domain data?

Prior work: Pretraining deep networks only on natural image statistics

Inspiration:
Learning to See By Looking at Noise (Baradad et al., NeurIPS 2021)

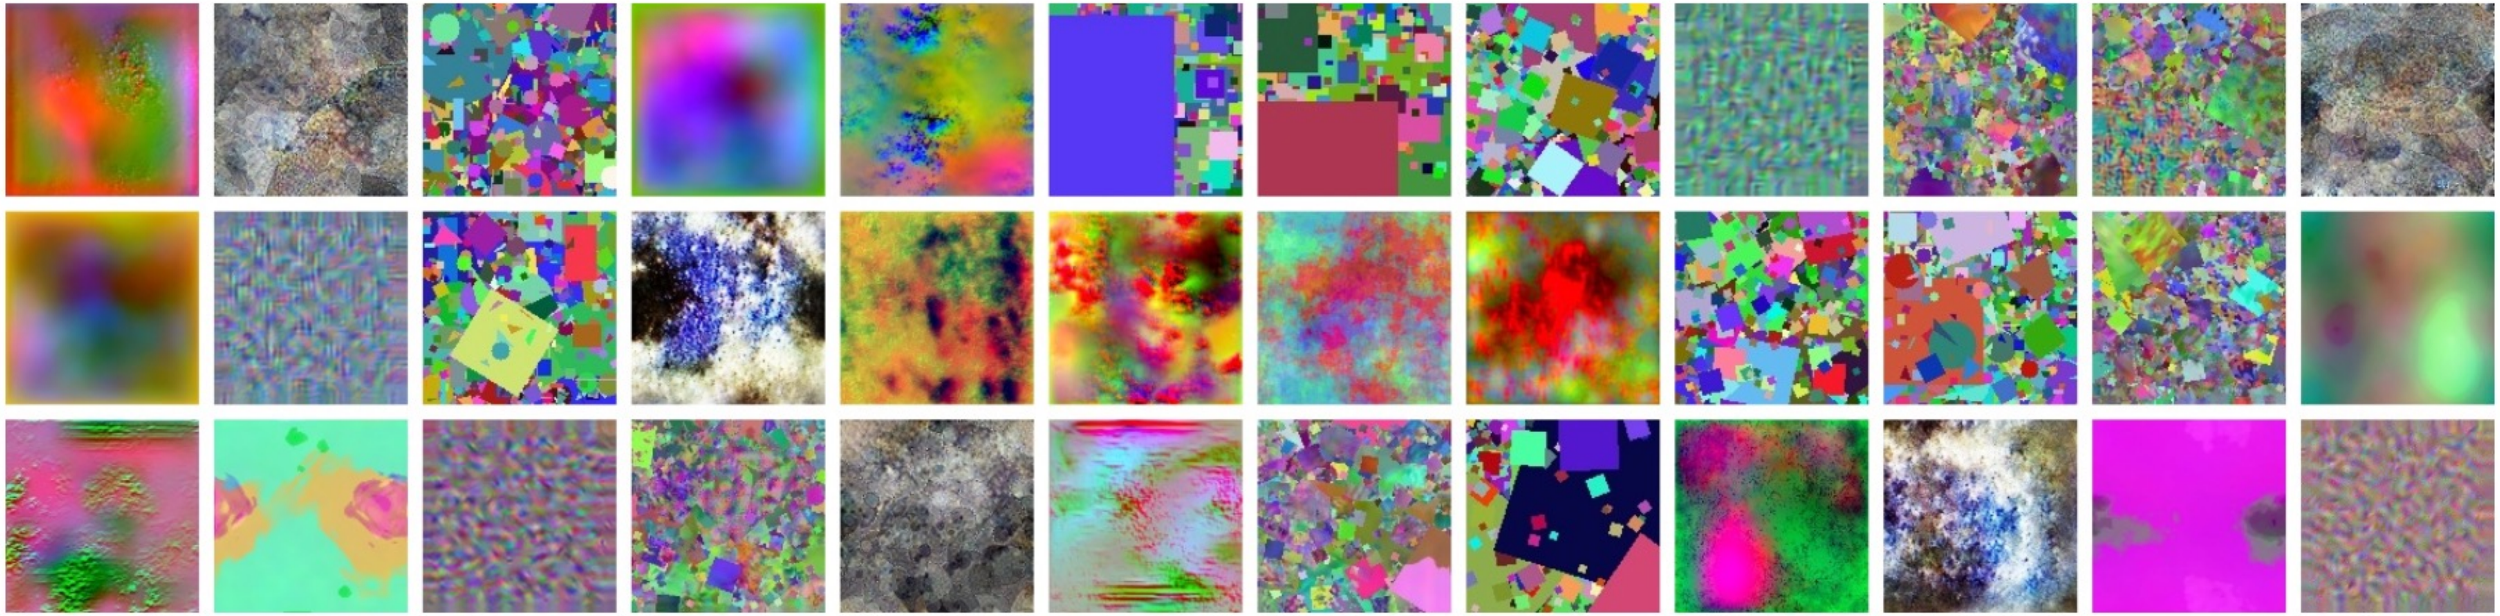

Main result: can pretrain visual classifiers using noise alone!

Type of noises used for pretraining SNP INRs

Structured Noise

Noise with $1/f^{\alpha}$ spectral structure

Dead-leaves noise

Captures natural statistics as prior

Unstructured Noise

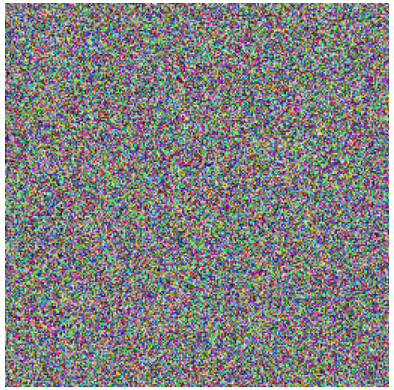

$\mathcal{U}(a, b)$

Forces INR to capture uncorrelated patterns

Structured noise image dataset from Baradad et.al.. *Learning to see by looking at Noise,* NeurIPS’21

SNP INRs: STRAINER with Noise Pretraining

Training: Fit on samples of a noise type

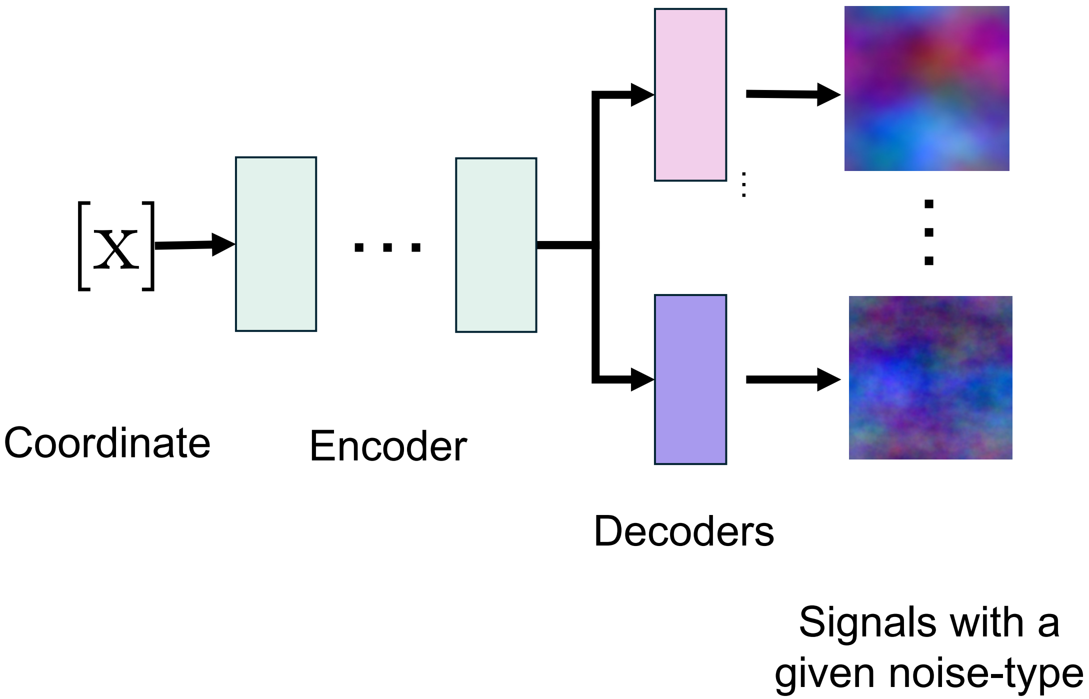

Testing: Fit on real signal

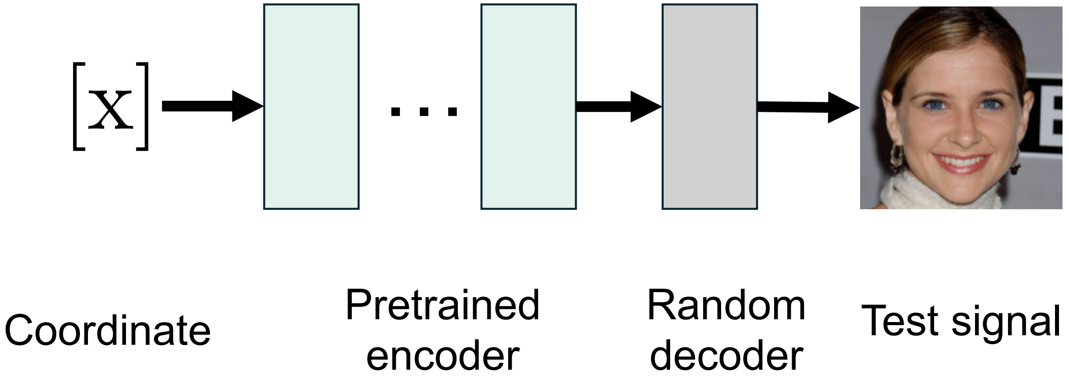

Results: SNP INRs are excellent signal fitters!

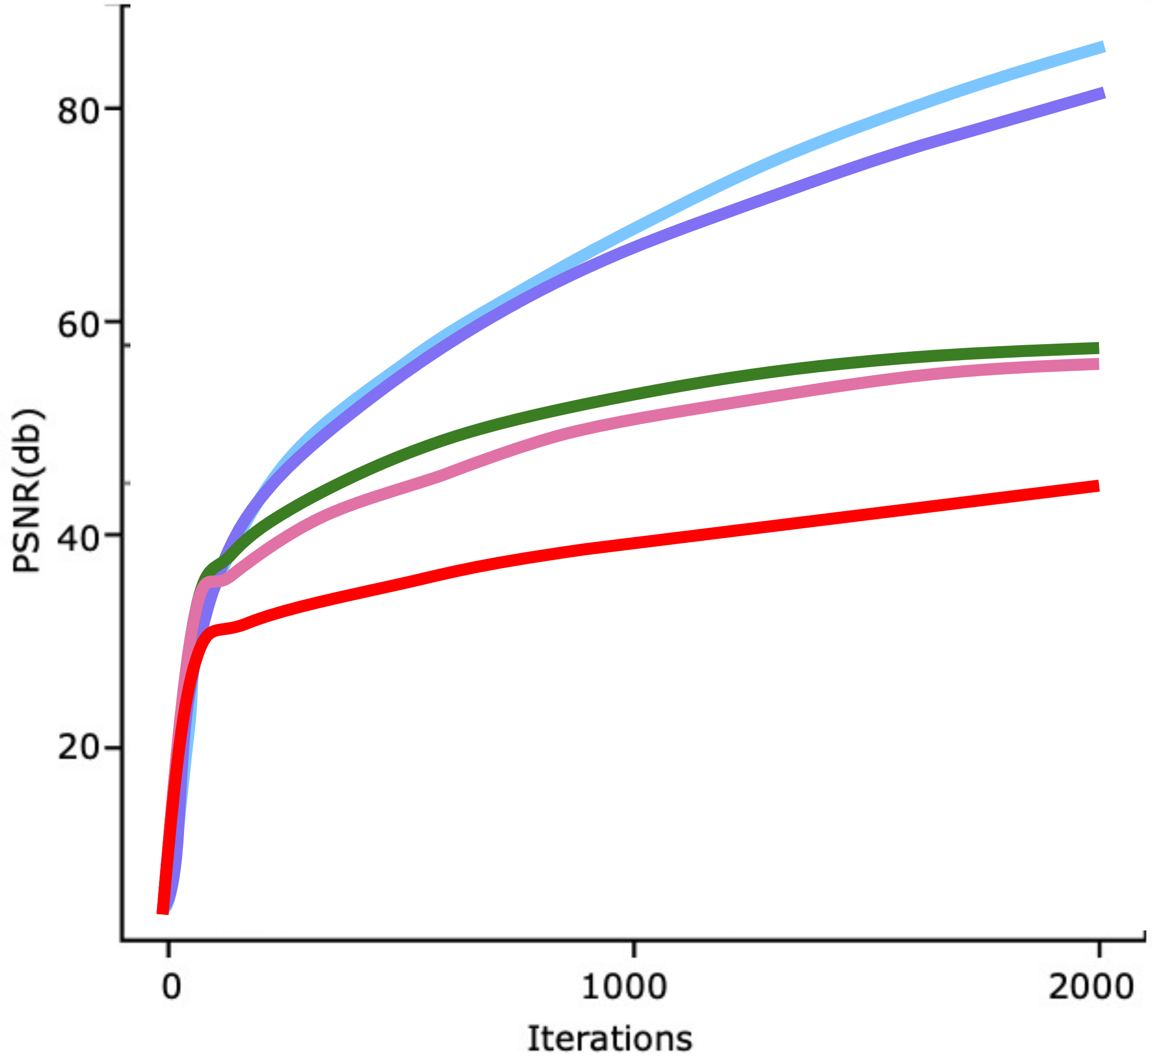

Results on 178 x 178 CelebA-HQ images

Baselines

- Random Init.
- STRAINER domain-data (CelebA-HQ)

SNP : pretrained on noise

- Uniform Noise
- Gaussian Noise
- 1/f Noise (Spectrum)

Frequency coverage during reconstruction

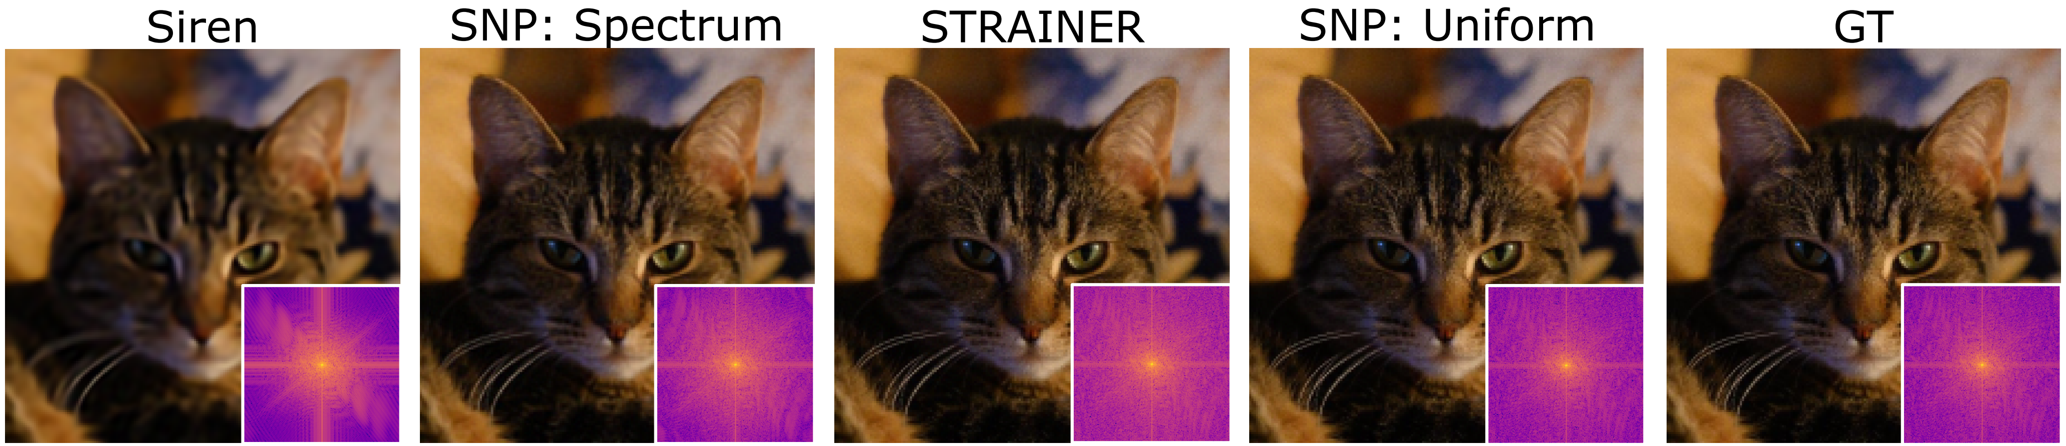

PSNR: 32.42, SSIM: 0.90, LPIPS: 0.21

PSNR: 39.17, SSIM: 0.98, LPIPS: 0.03

PSNR: 46.42, SSIM: 0.99, LPIPS: 0.003

PSNR: 53.34, SSIM: 0.99, LPIPS: 0.0003

Frequency coverage during reconstruction

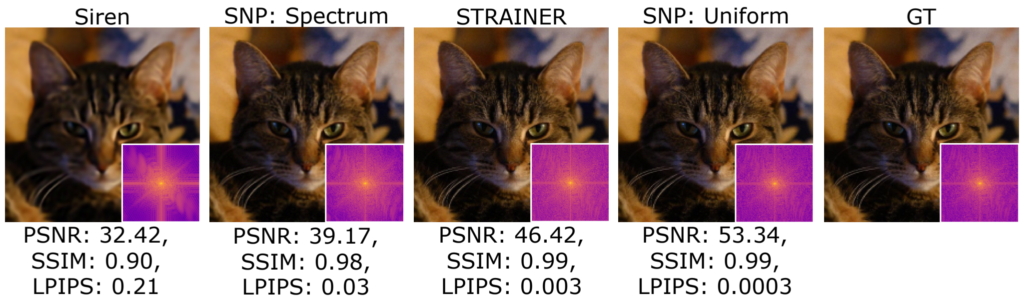

Neural tangent kernel (NTK) analysis

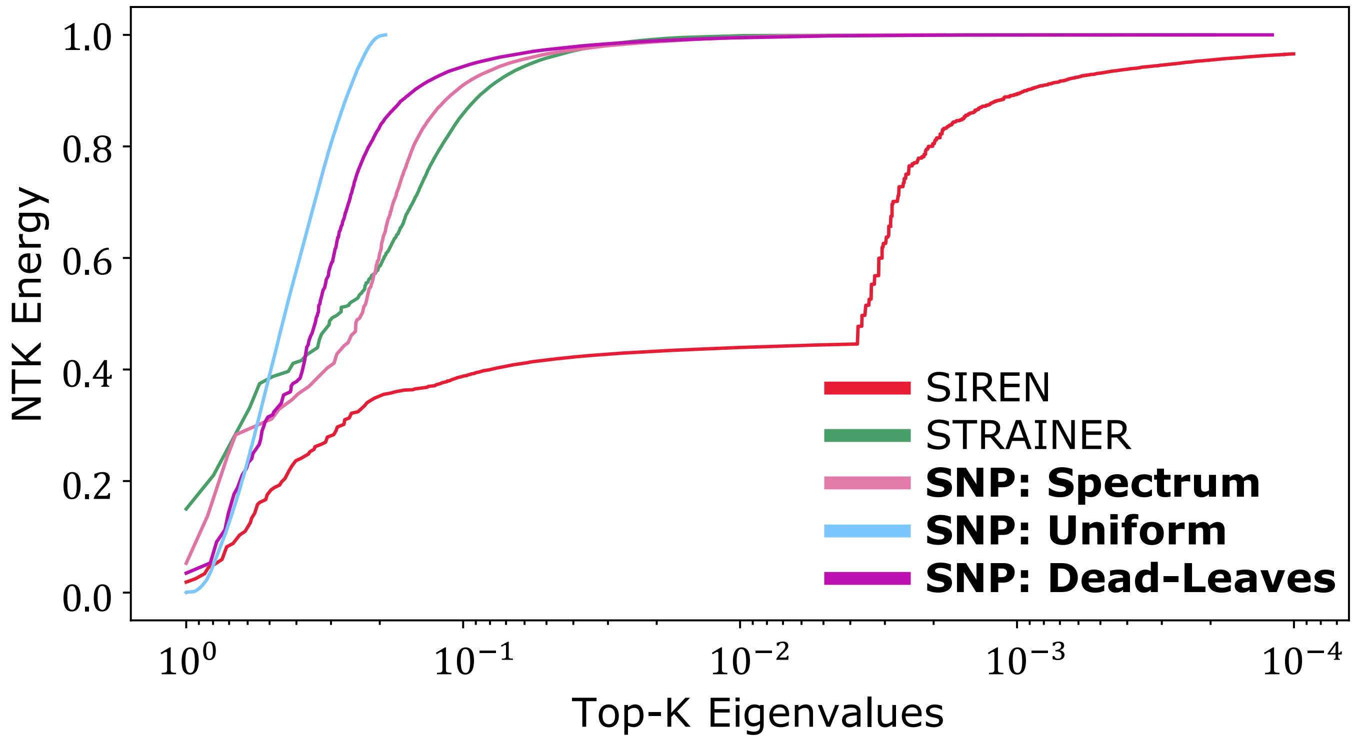

SNP INRs readily extend to video fitting and denoising

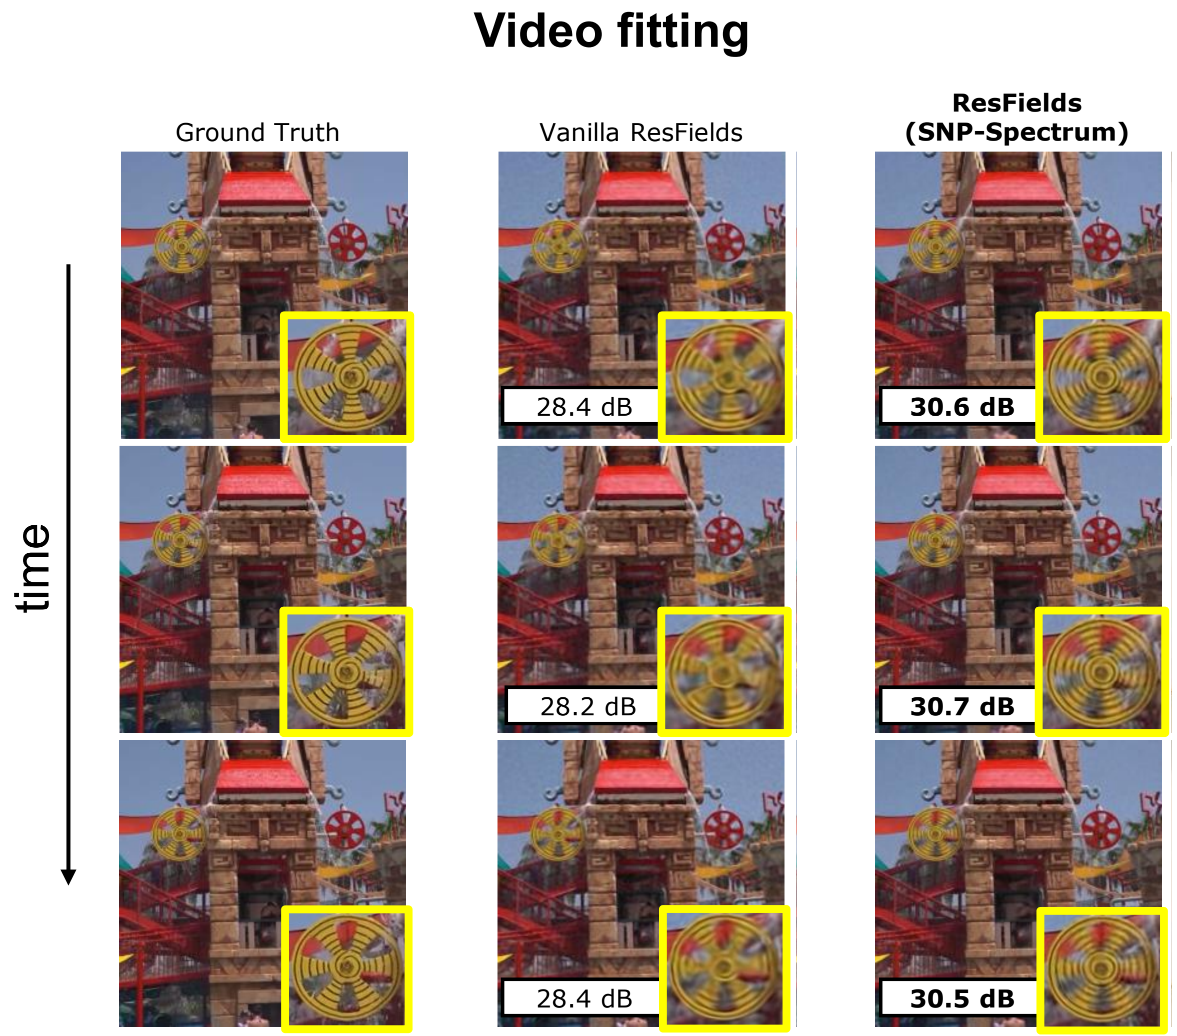

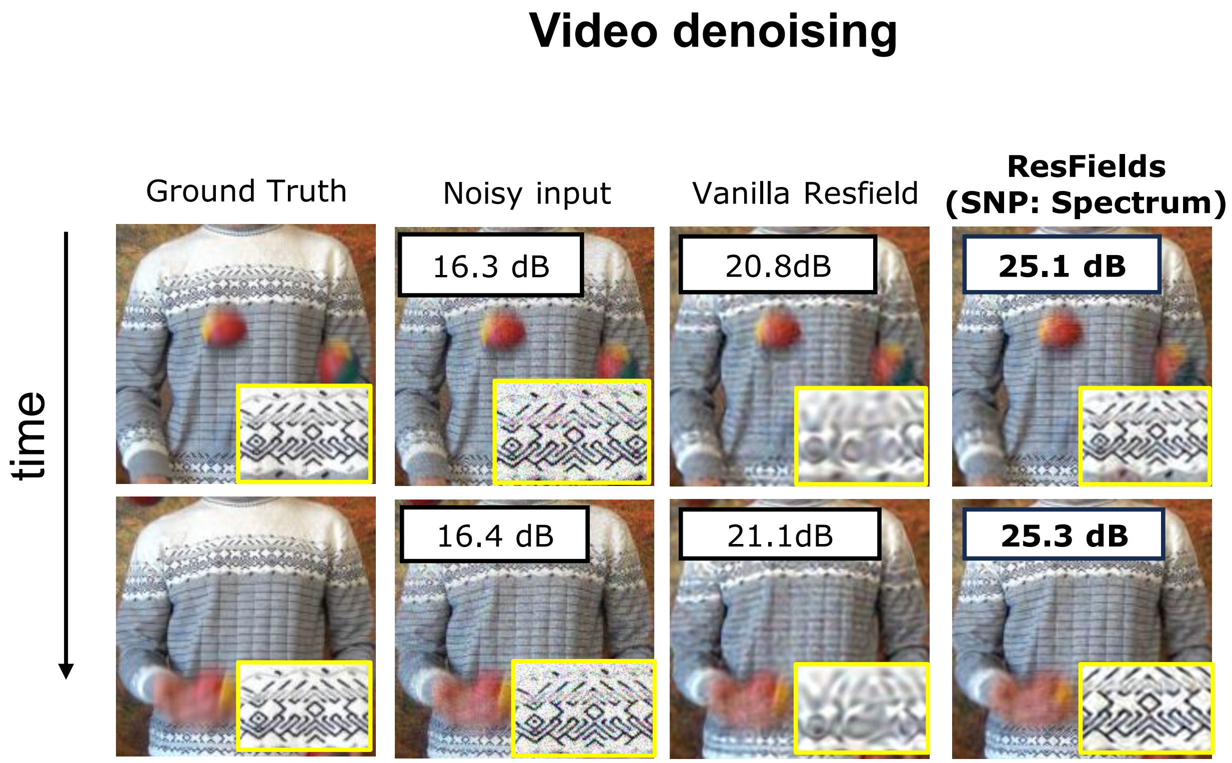

Denoising vs. Fitting (CelebA-HQ)

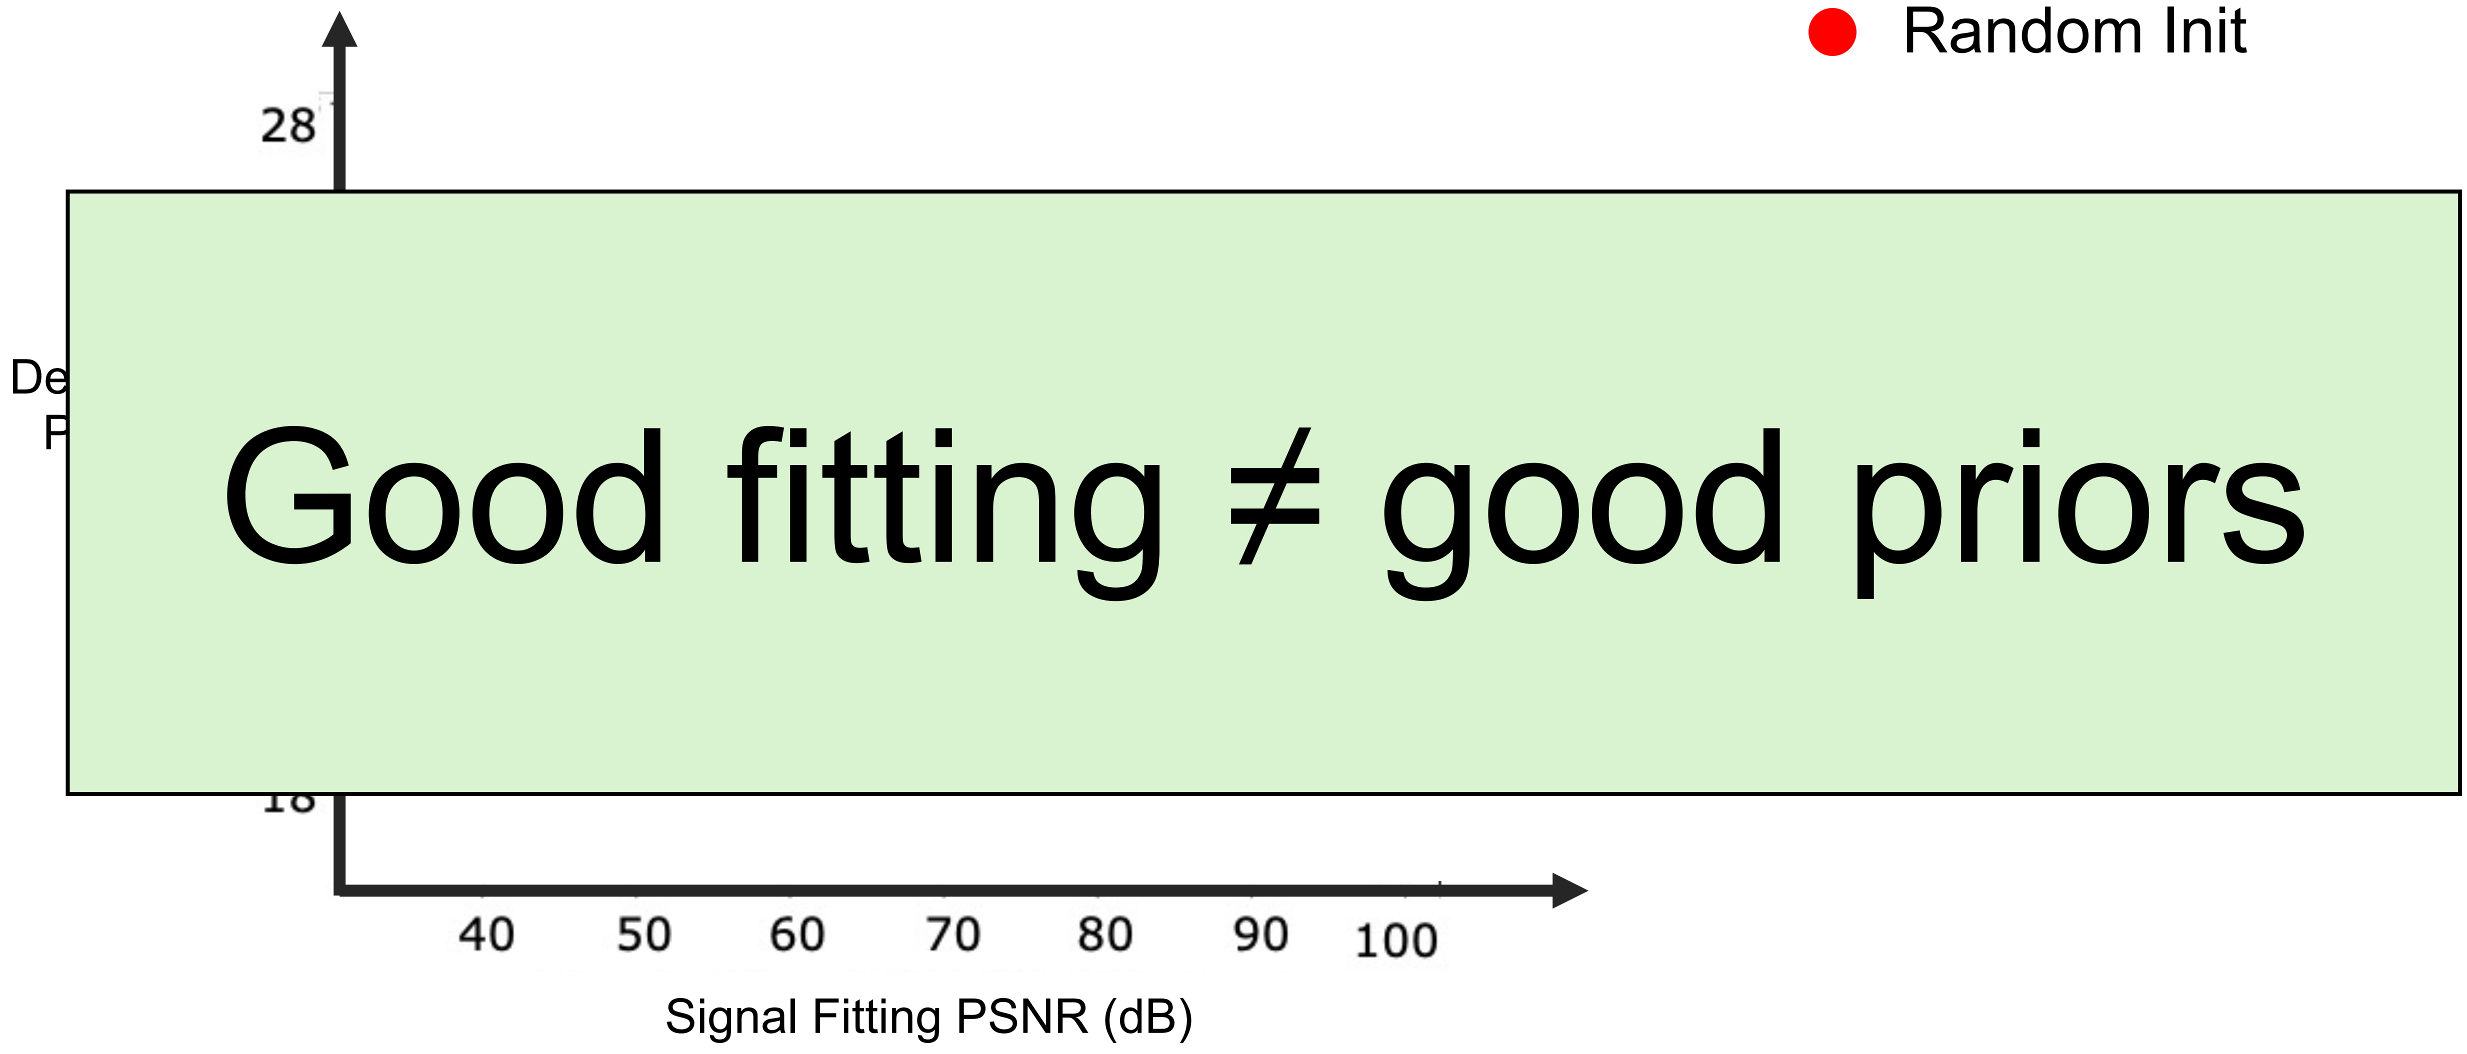

Denoising vs. Fitting (CelebA-HQ)

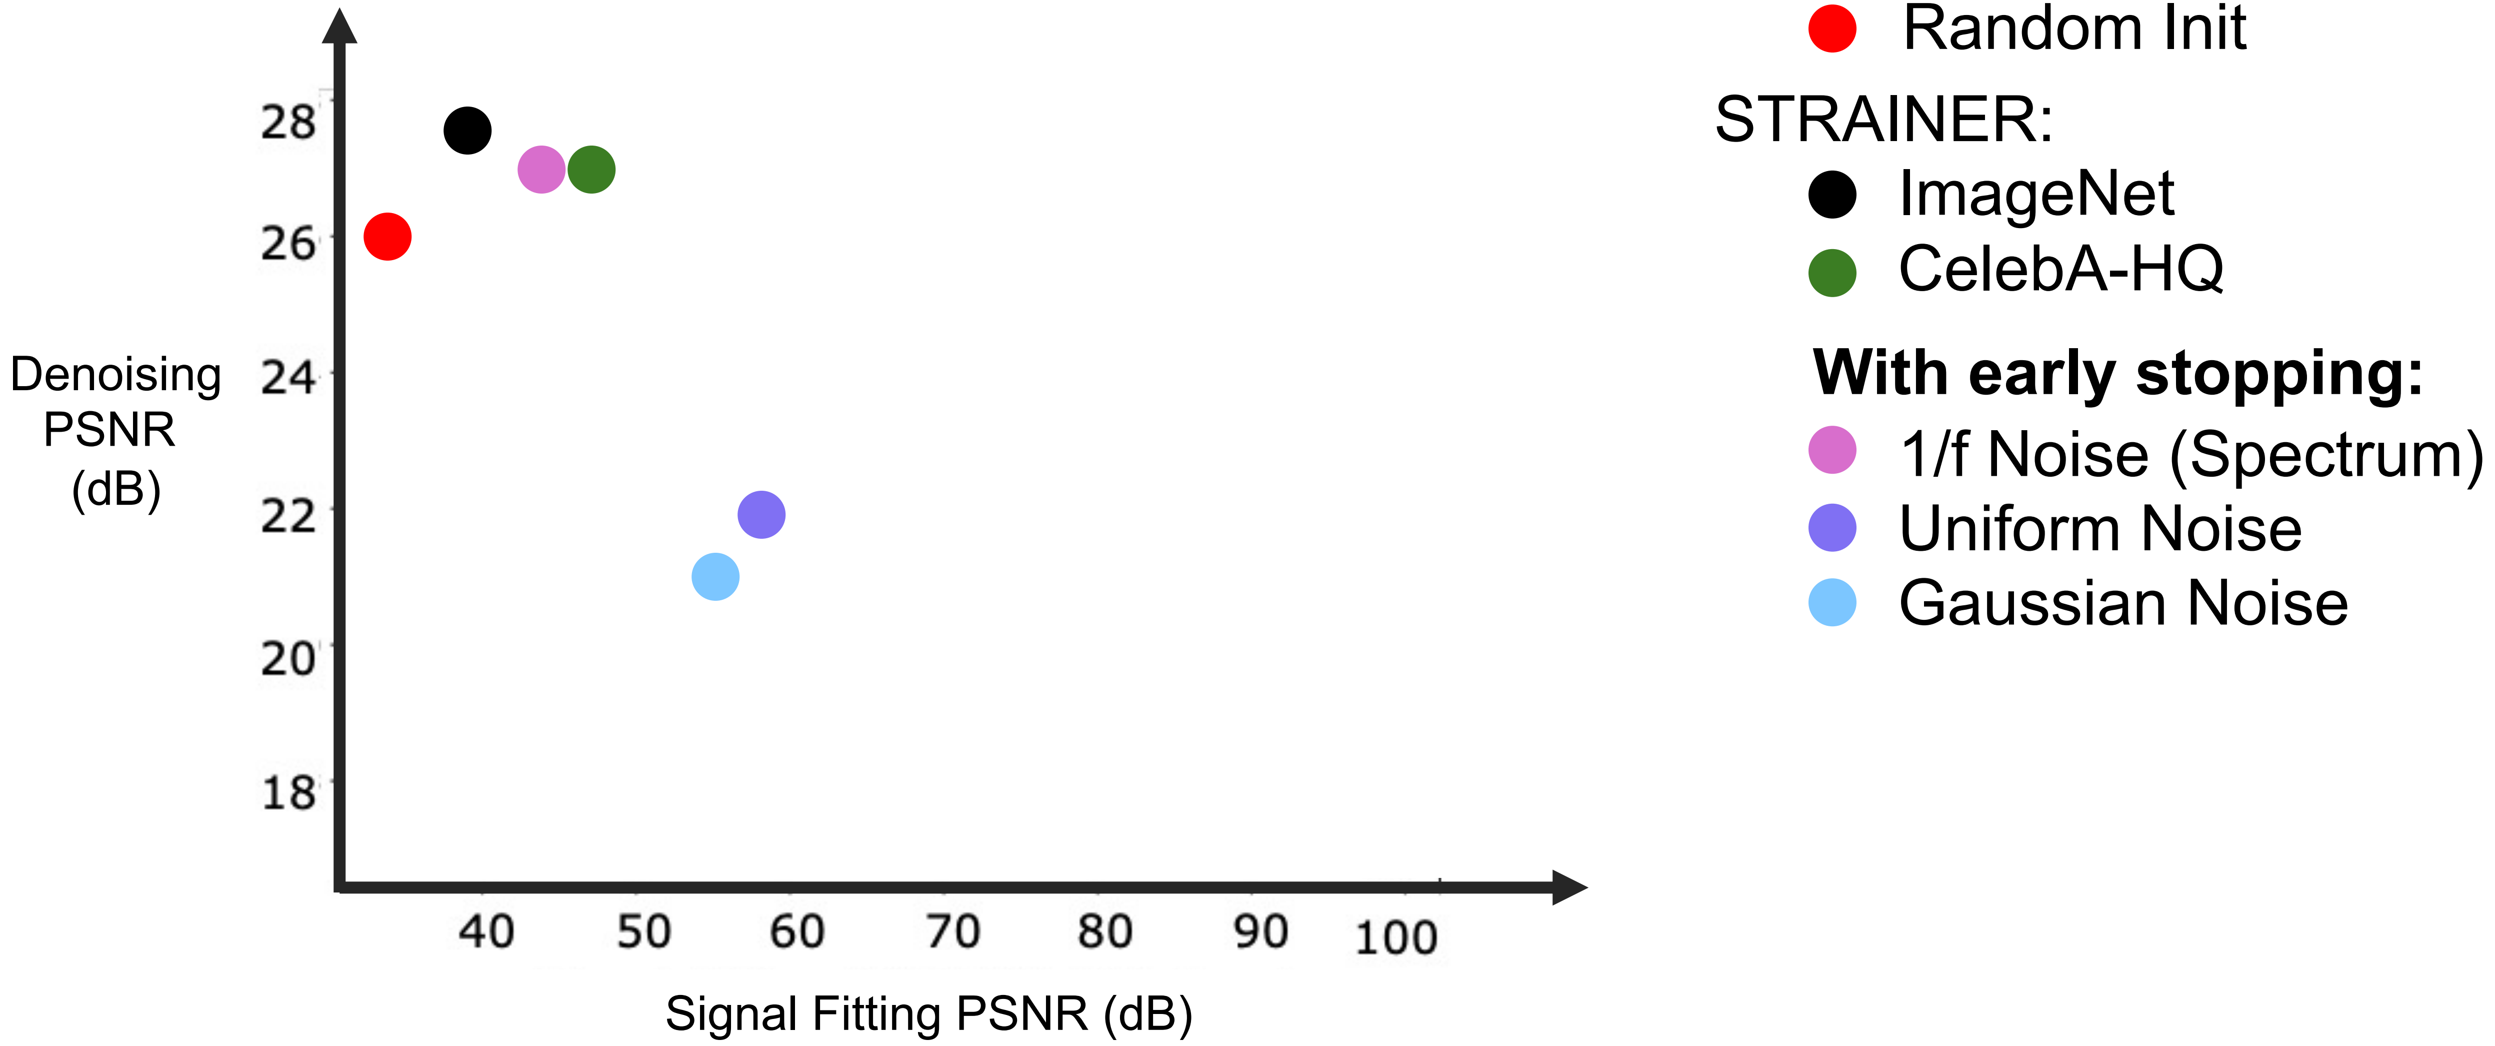

Denoising vs. Fitting (CelebA-HQ)

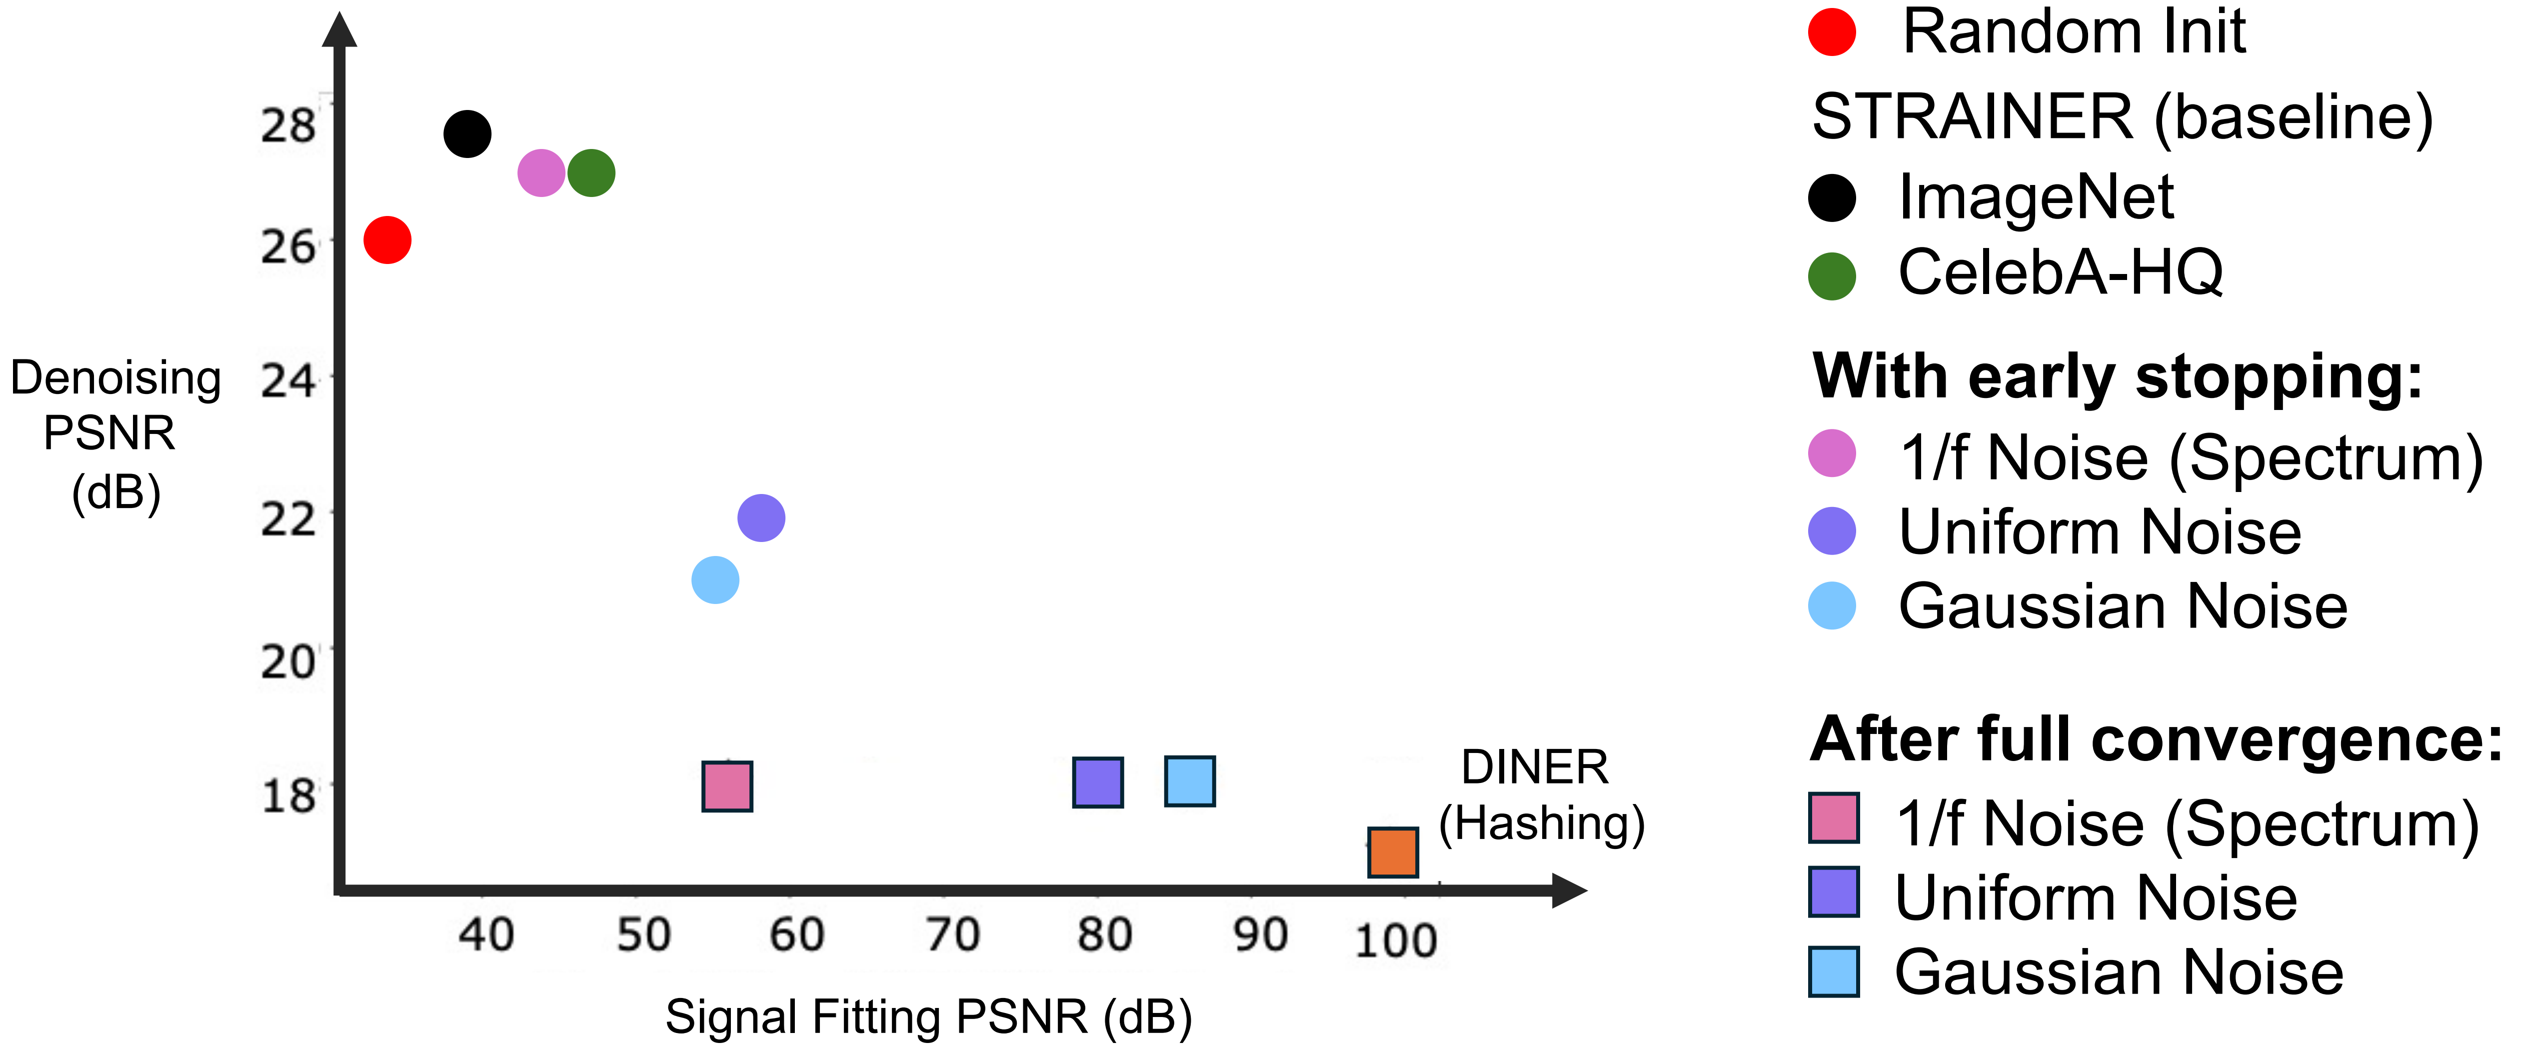

What does pretraining encode in the INR?

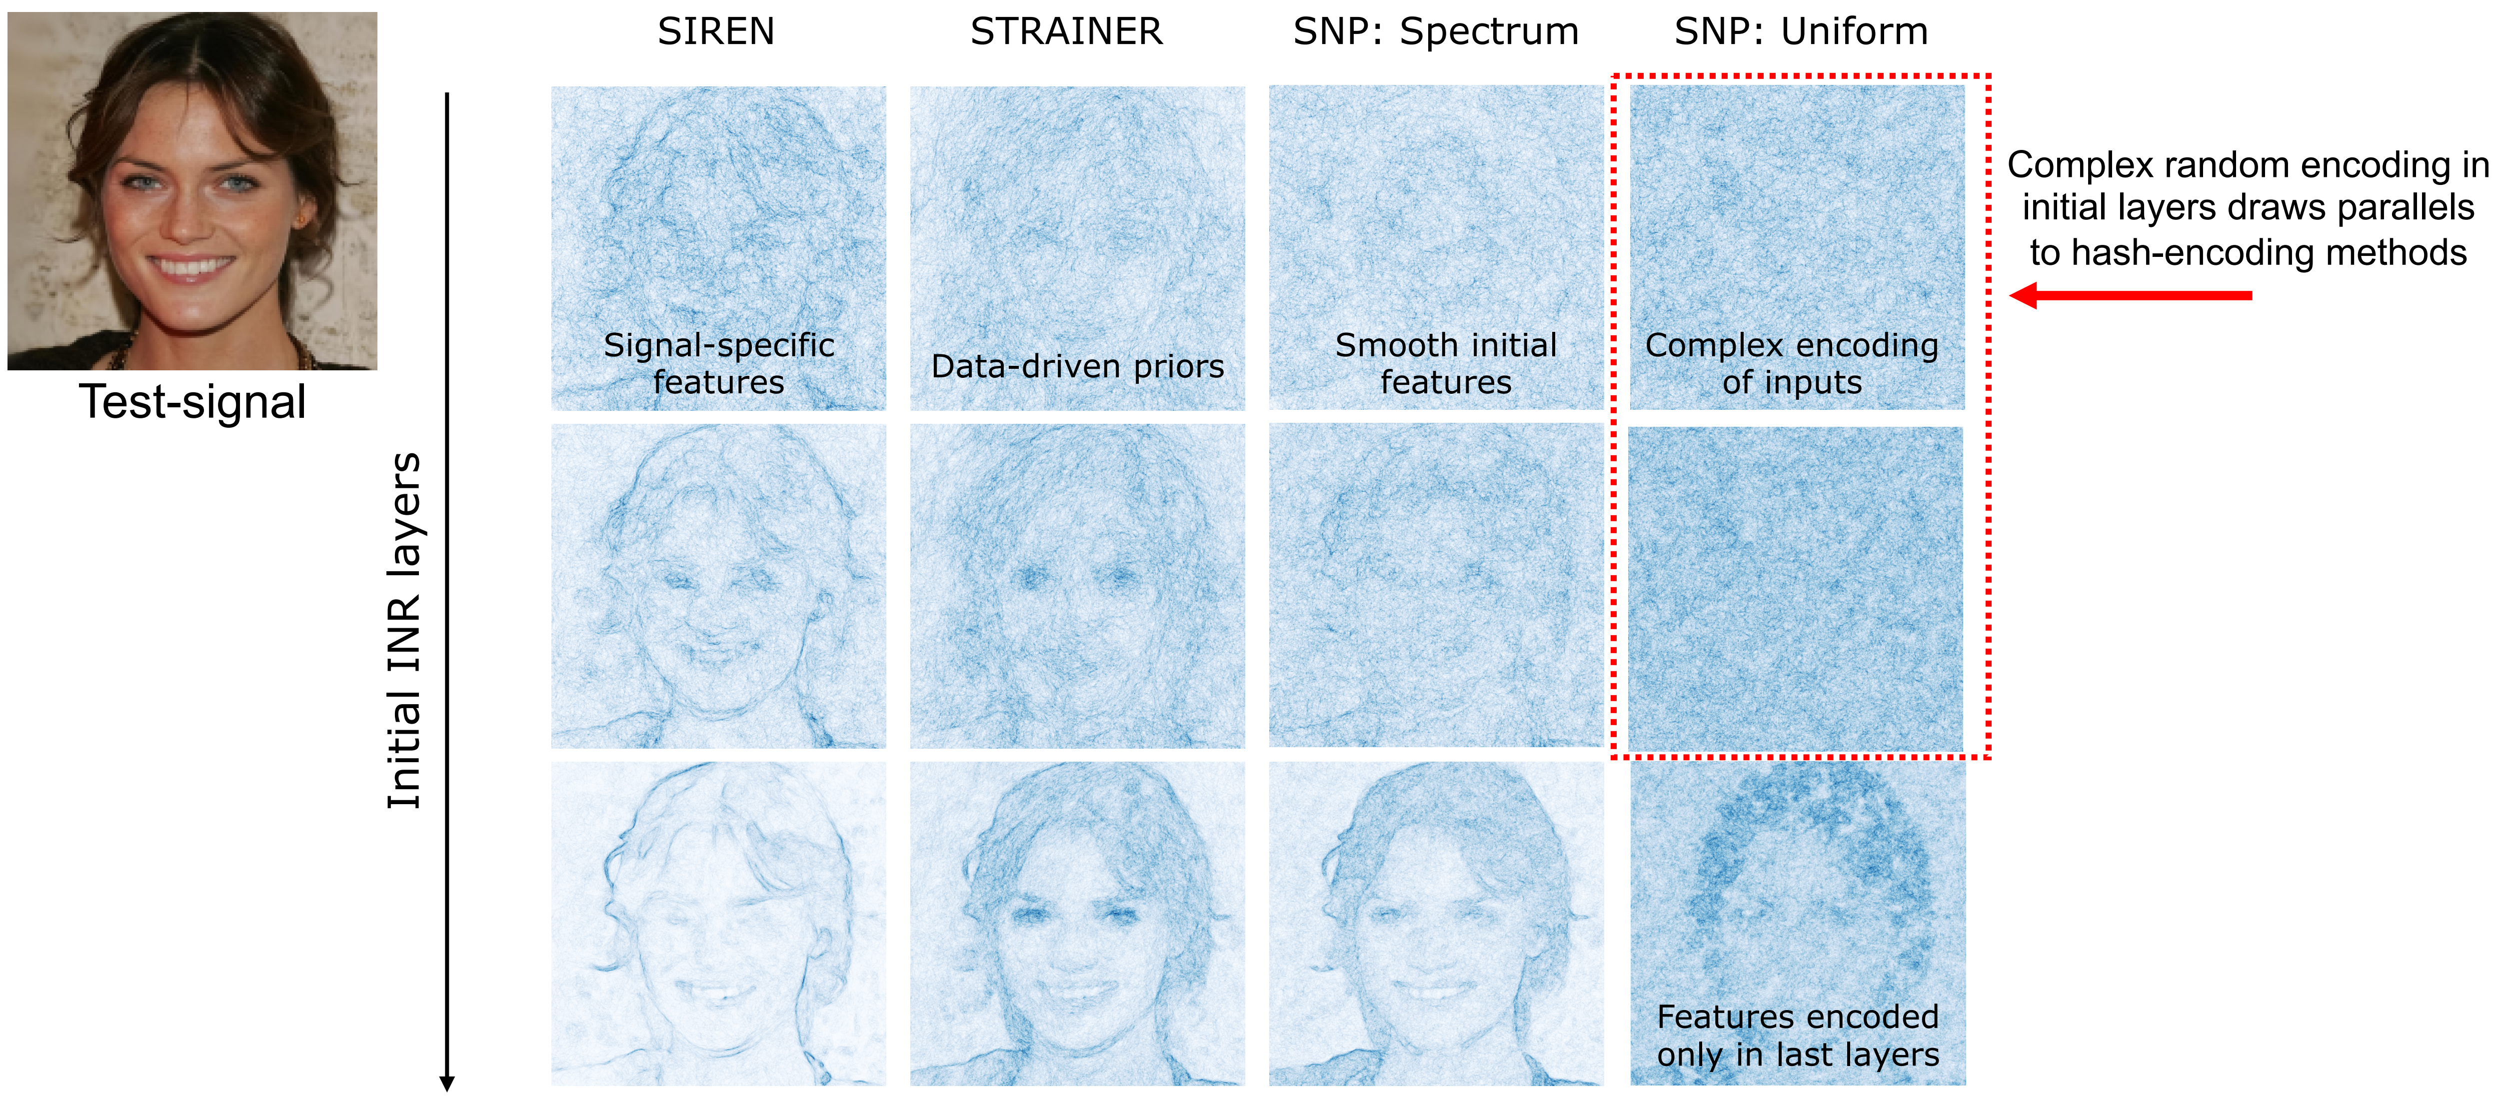

Visualizing the loss landscape of SNP INRs

SIREN | STRAINER | SNP: Spectrum | SNP: Uniform

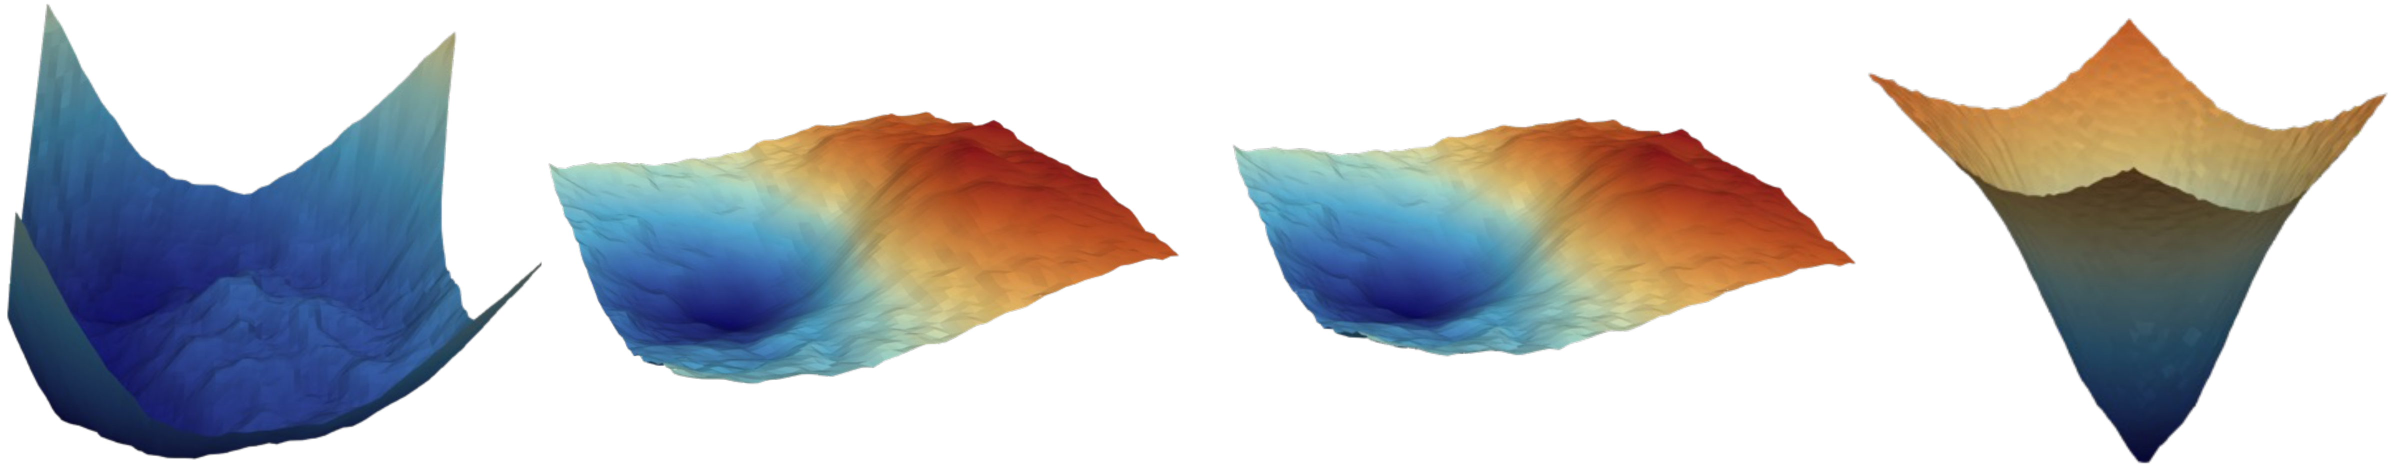

Rugged landscape, unclear minimum

STRAINER and SNP Spectrum have near-identical loss landscapes

Smooth landscape, distinctive minima

Summary and contributions

- We propose SNP INRs : a novel data-free initialization technique for implicit neural representations (INRs) using structured and unstructured noises.
- An overall recommendation of our study is to pretrain INRs on noise with $1/f^{\alpha}$ spectral structure to obtain fitting quality and inverse problem (denoising) performance at par with data-driven methods.
- We also show the surprising finding that pretraining on unstructured noise such as Uniform and Gaussian yields excellent high-quality signal fitters, with behavior similar to hash-encoding based methods.
- We support our study with more detailed analysis using neural tangent kernel, visualizing the loss landscape of SNP INRs, and

SNP INRs: The Surprising Effectiveness of Noise Pretraining for Implicit Neural Representations

Kushal Vyas[1], Alper Kayabasi[2], Daniel Kim[1], Vishwanath Saragadam[2], Ashok Veeraraghavan[1], Guha Balakrishnan[1]

[1]Rice University, [2]University of California, Riverside


Kushal
Vyas


Alper
Kayabasi


Daniel
Kim


Vishwanath
Saragadam


Ashok
Veeraraghavan


Guha
Balakrishnan

Poster #567. Poster session #1.
ExHall A-F, June 5, 2026.

Project Page