NEURIPS 2026 · MAIN CONFERENCE

Closed-Form Implicit Neural Representations

Woojin Cho **Junghwan Park** TelePIX woojin@telepix.net · junghwan@telepix.net

PAPER
arXiv:2602.07603

TELEPIX
EXTRACTING INSIGHTS FROM SPACE

01 INRs are slow to fit

An INR stores a signal as a continuous function of its coordinates. Each one must still be fitted from scratch.

Iterative backpropagation. Thousands of gradient steps per signal: slow in wall-clock time, nonconvex, hard to interpret.

Spectral bias. Networks learn low frequencies first; fine-scale detail arrives late, or is never fitted.

02 ELM-INR: solve, don't train

STEP 1 **Decompose the domain**

Split the domain Ω into N overlapping subdomains Ω_i.

STEP 2 **Fit each subdomain in closed form**

A local Extreme Learning Machine keeps its random hidden layer frozen. Only the output weights are learned, which is a linear least-squares problem.

LOCAL ELM
$$\hat{f}_i(\mathbf{x}) = \Sigma_j \, \alpha_{ij} \, \sigma(\mathbf{w}_{ij}{}^{\mathrm{T}}\mathbf{x} + b_{ij})$$

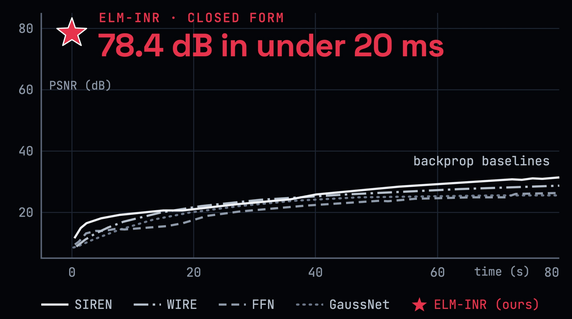
CLOSED FORM
$$\alpha_i{}^{*} = (\mathbf{H}_i{}^{\mathrm{T}}\mathbf{H}_i)^{-1}\,\mathbf{H}_i{}^{\mathrm{T}}\,\mathbf{y}_i$$

STEP 3 **Blend with a partition of unity**

Smooth windows φ_i stitch the local fits into one continuous signal. The solves stay decoupled, so they run in parallel on a GPU.

PARTITION OF UNITY
$$\hat{f}(\mathbf{x}) = \Sigma_i \, \varphi_i(\mathbf{x})\, \hat{f}_i(\mathbf{x}), \quad \Sigma_i \, \varphi_i(\mathbf{x}) = 1$$

03 Milliseconds, not minutes

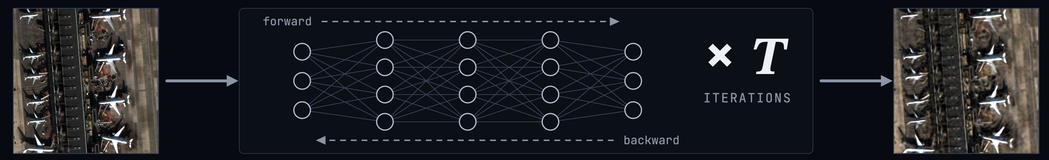


kodim05, 512 × 512 RGB. After 1,000 epochs the strongest baseline (SIREN) reaches 30.5 dB. ELM-INR reaches 78.4 dB without a single gradient step.

CONCEPT

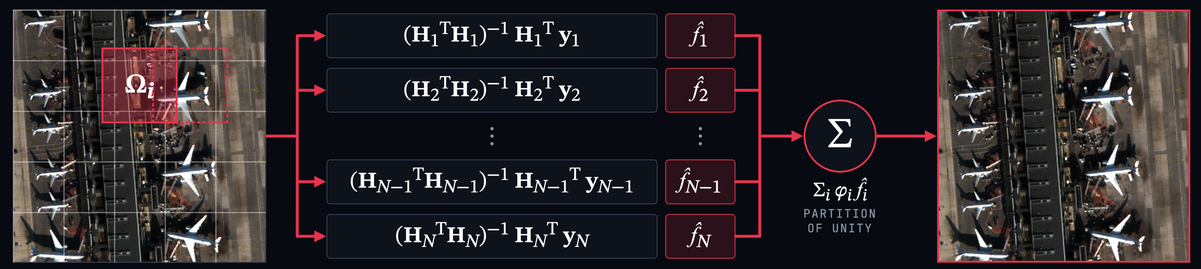
Swap the training loop for one linear solve per subdomain.

BACKPROPAGATION-BASED INR



Input Ω — Iterative update: forward and backward passes, repeated T times — Output $\hat{f}$

BACKPROPAGATION-FREE ELM-INR · OURS

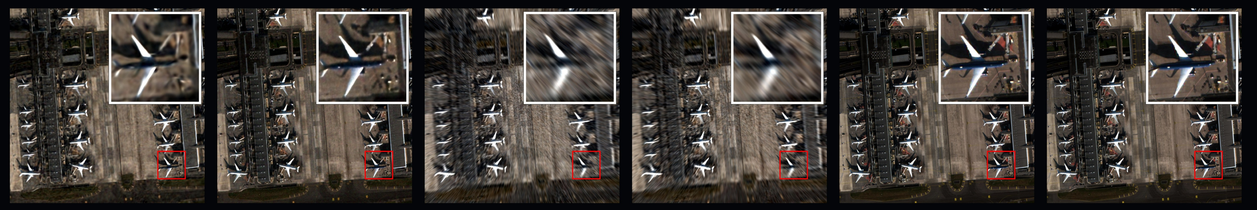


Overlapping subdomains $\Omega_1 \ldots \Omega_N$ — N independent least-squares solves, batched on a GPU

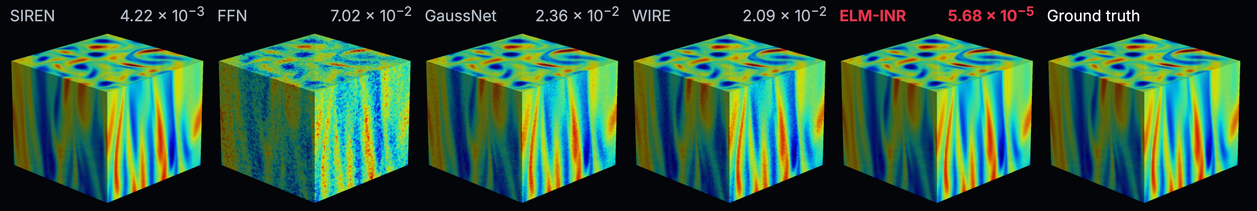
SCENE: WORLDVIEW-3 MULTISPECTRAL, PARIS-CDG AIRPORT · 8 BANDS · 2048 × 2048 · RGB CROP SHOWN

04 Near-exact reconstruction across modalities

WORLDVIEW-3 MULTISPECTRAL · 8 BANDS · 2048 × 2048 · PSNR, HIGHER IS BETTER



NAVIER-STOKES VOLUME · SPACE × SPACE × TIME · MAE, LOWER IS BETTER



05 Where is fitting hard?

The **spectral Barron norm** weighs a signal's frequencies by their magnitude: a proxy for how hard it is to fit.

BARRON NORM
$$\|f\|_{B_S} = \int \|\xi\|_1 \, |\mathcal{F}(\xi)| \, \mathrm{d}\xi$$

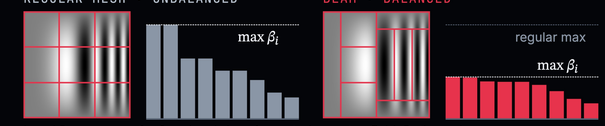
With m hidden units per subdomain, the worst local error is set by the most complex region:

ERROR BOUND
$$\max_i \|f - \hat{f}_i\|_{L^2(\Omega_i)} \leq \frac{C}{\sqrt{m}} \cdot \max_i \beta_i$$

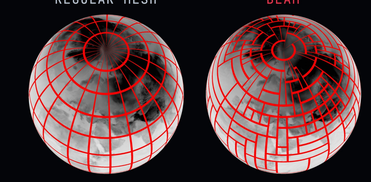
$\beta_i = \|f|_{\Omega_i}\|$ is the local Barron norm on subdomain Ω_i.

06 BEAM balances the spectrum

Barron-Enhanced Adaptive Mesh merges atomic cells bottom-up while the merged norm stays below a threshold τ: large patches where the signal is smooth, small ones where it is complex.



Local Barron norm β_i of each subdomain, sorted, on a toy field. Both meshes use nine subdomains.



SAME N, SAME m

ERA5 temperature · MAE
m = 1024 0.398 → **0.363**
m = 2048 0.317 → **0.288**

kodim20 · PSNR
m = 256 26.0 → **27.5 dB**
m = 512 29.6 → **32.3 dB**

07 In orbit: ELMZip 2.0

SPOTLIGHT · TACKLING CLIMATE CHANGE WITH ML WORKSHOP, NEURIPS 2026

The same closed-form fit, turned into an onboard codec. A patch-invariant basis reduces encoding to **one matrix multiplication**: no training and no solver on the satellite.

~20× compression on captures taken in orbit

1 GEMM is the whole onboard coefficient fit

15 W Jetson-class flight computer on BlueBON



BlueBON 6U CubeSat capture of Bordeaux, reconstructed from the payload encoded in orbit.

THIS WORK WAS SUPPORTED BY TELEPIX