UNIVERSITY OF MARYLAND

acm SIGSPATIAL 34th ACM SIGSPATIAL Riverside, CA, 2026

POSTER R-110

Rethinking Amortized Neural Representations for High-Resolution Terrain Elevation Data

University of Maryland, College Park

Haoan Feng, Xin Xu, Leila De Floriani

Key contributions

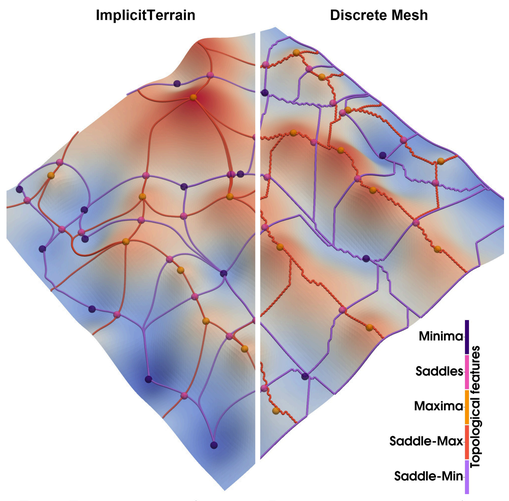
Terrain benchmark for INRs
Compare TransINR, Functa and HUVR on high-resolution terrain DEMs.

Adaptive neural decoder for terrain
Improve elevation and derivative fidelity without enlarging the tile embedding.

Lightning conversion and compact storage
Millisecond feed-forward encoding time for each tile 8-bit compact embedding via shared models.

01 INRs for terrain modeling and analysis

ImplicitTerrain, our prior work at CVPR Workshops 2024[1]



Continuous queries
Elevation at arbitrary coordinates

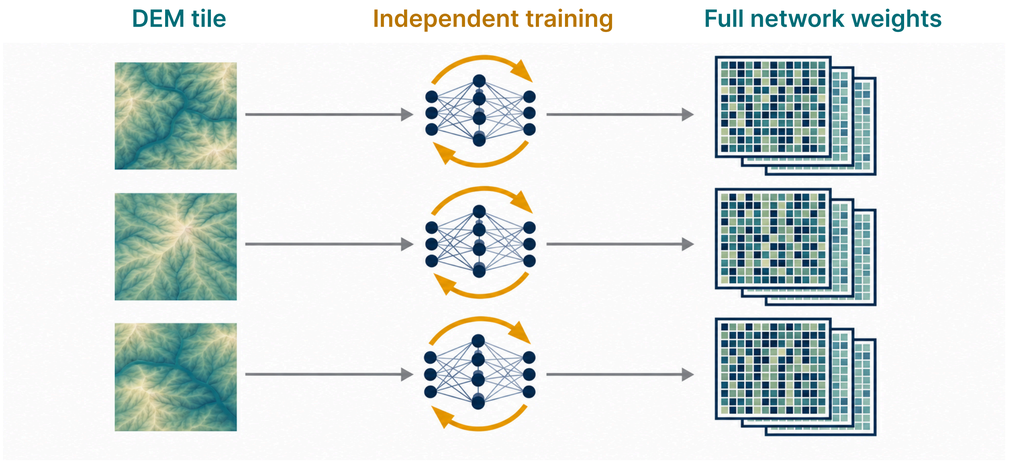
Flexible sampling
Evaluate one model at different sampling resolutions

Analytic derivatives
Slope, aspect and curvature via analytical coordinate derivatives

Terrain analysis
Critical features and topological structure

02 Per-instance fitting: time and storage

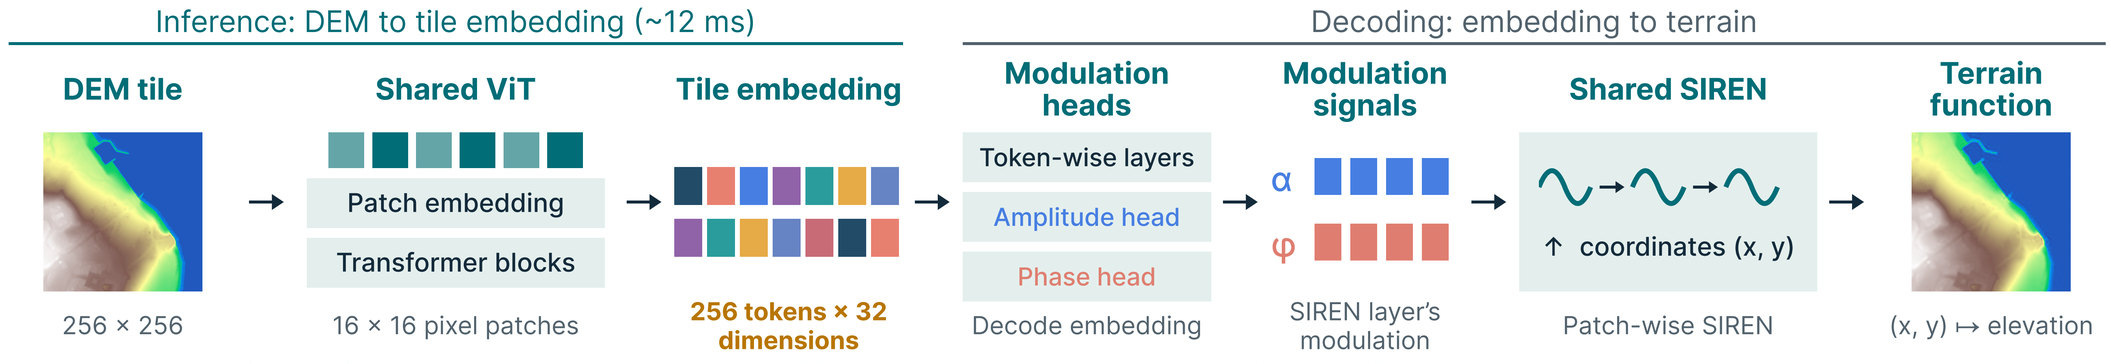
Finer terrain detail requires more parameters in each independently trained and stored network.



Training cost ∝ N; Storage cost ∝ N

03 Amortized conversion with a terrain-aware decoder

Benchmark TransINR[2], Functa[3] and HUVR[4]. Our **HUVR+SIREN** adaptation replaces HUVR's ReLU-MLP-based decoder and up-sampler with a modulated SIREN[5], representing terrain elevation as a differentiable coordinate function.



End-to-end training of shared models on our curated DEM dataset; store one embedding per tile.

Encode a tile, then query its terrain function

$$\zeta^{(i)} = T_\Phi(S^{(i)}), \qquad \hat{s}^{(i)}(\mathbf{x}) = f(\mathbf{x}; \bar{\theta}, \zeta^{(i)})$$

$S^{(i)}$ input DEM tile *i*
$\zeta^{(i)}$ tile embedding predicted by the encoder
$\mathbf{x},\ \hat{s}^{(i)}(\mathbf{x})$ query coordinate and predicted elevation for terrain function *f*
$\Phi,\ \bar{\theta}$ encoder and decoder weights shared across tiles

Amplitude- and phase-modulated SIREN

$$\mathbf{h}_\ell = \boldsymbol{\alpha}_\ell^{(p)} \odot \sin\left(\omega_0(W_\ell \mathbf{h}_{\ell-1} + \mathbf{b}_\ell) + \boldsymbol{\varphi}_\ell^{(p)}\right)$$

$\boldsymbol{\alpha}_\ell^{(p)},\ \boldsymbol{\varphi}_\ell^{(p)}$ amplitude and phase predicted for patch *p*
$\mathbf{h}_{\ell-1},\ \mathbf{h}_\ell$ input and output of SIREN layer *l*
$W_\ell,\ \mathbf{b}_\ell$ SIREN layer: shared layer weights and bias
$\omega_0,\ \odot$ frequency scale and element-wise product

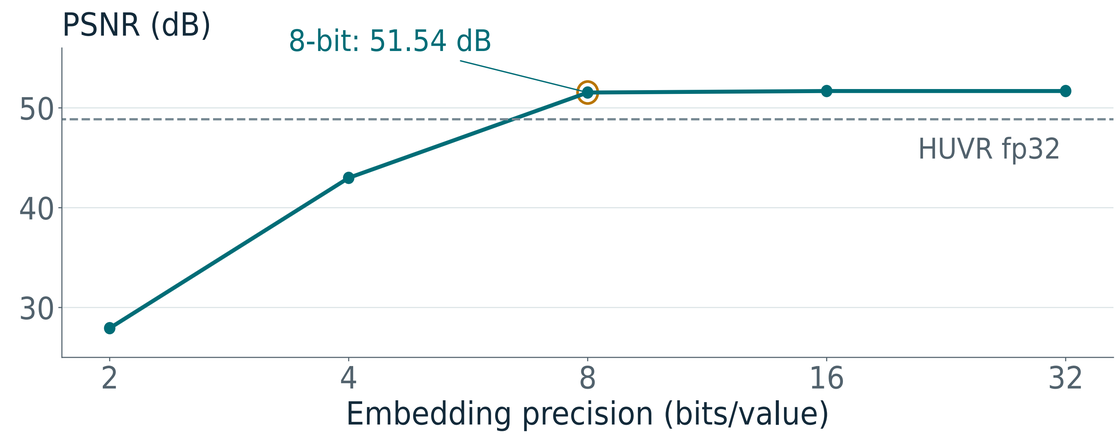
04 Terrain fidelity and embedding storage

swissALTI3D, swisstopo[6] 432 test tiles, 256 × 256 pixels, evaluated at 1 m/pixel

Method	PSNR (dB) ↑	RMSE (m) ↓
TransINR[2]	41.56	0.222
Functa[3]	35.00	0.498
HUVR[4]	48.86	0.092
HUVR+SIREN (ours)	**51.69**	**0.066**

Compared with HUVR: +2.83 dB PSNR and lower elevation RMSE (0.092 → 0.066 m), at the same embedding size.

Embedding precision and reconstruction quality

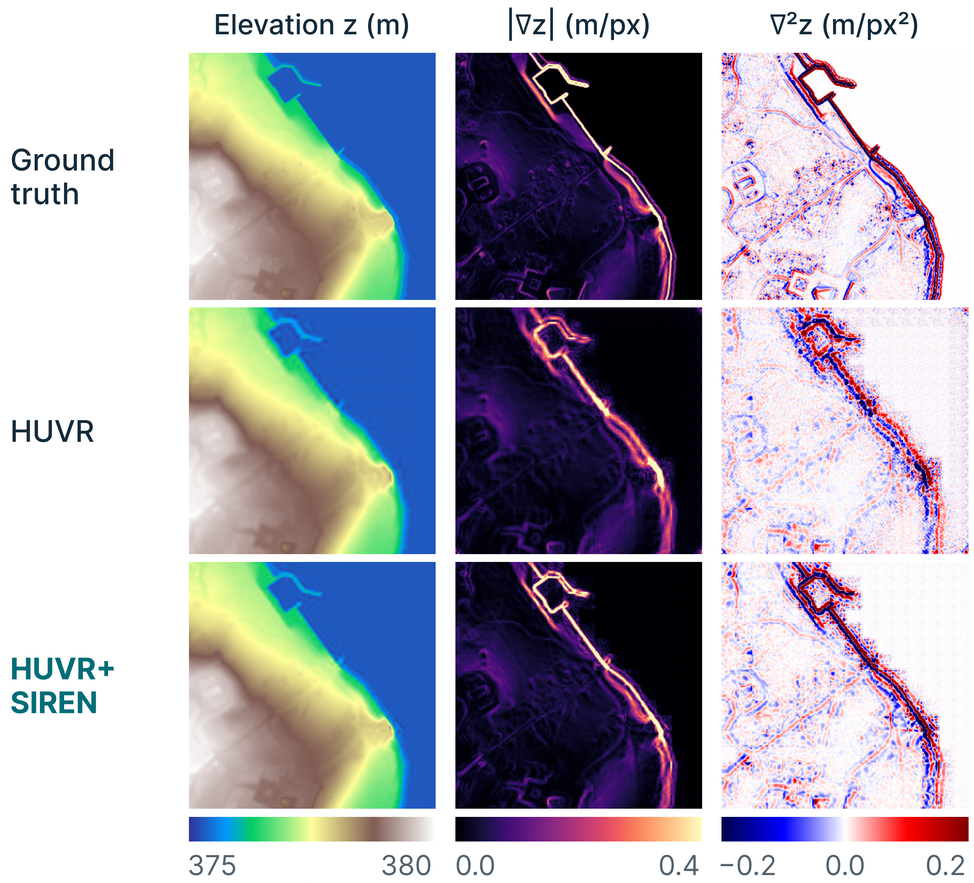


8-bit embeddings retain 51.54 dB PSNR: a 0.15 dB loss from fp32, with ×4 fewer bits per stored value.

DEM-to-embedding feed-forward inference: 11.76 ms per tile on RTX A5000.

05 Terrain geometry and derivative fidelity

Qualitative comparison on a held-out test tile.



SIREN decoder better retains fine structure in the derivative maps.
Gradient magnitude characterizes slope; the Laplacian captures second-order elevation variation. Across 432 held-out tiles, their RMSE falls by 23% and 17%, respectively, compared with HUVR.

References

[1] Feng et al. ImplicitTerrain. CVPRW 2024.
[2] Chen & Wang. TransINR. ECCV 2022.
[3] Dupont et al. Functa. ICML 2022.
[4] Gwilliam et al. HUVR. arXiv:2601.14256, 2026.
[5] Sitzmann et al. SIREN. NeurIPS 2020.
[6] Federal Office of Topography swisstopo. swissALTI3D.

Haoan Feng
Ph.D. Candidate, Computer Science
University of Maryland, College Park

Actively seeking full-time opportunities

fengyee.github.io
hfengac@umd.edu

Project

Homepage