Quillwave

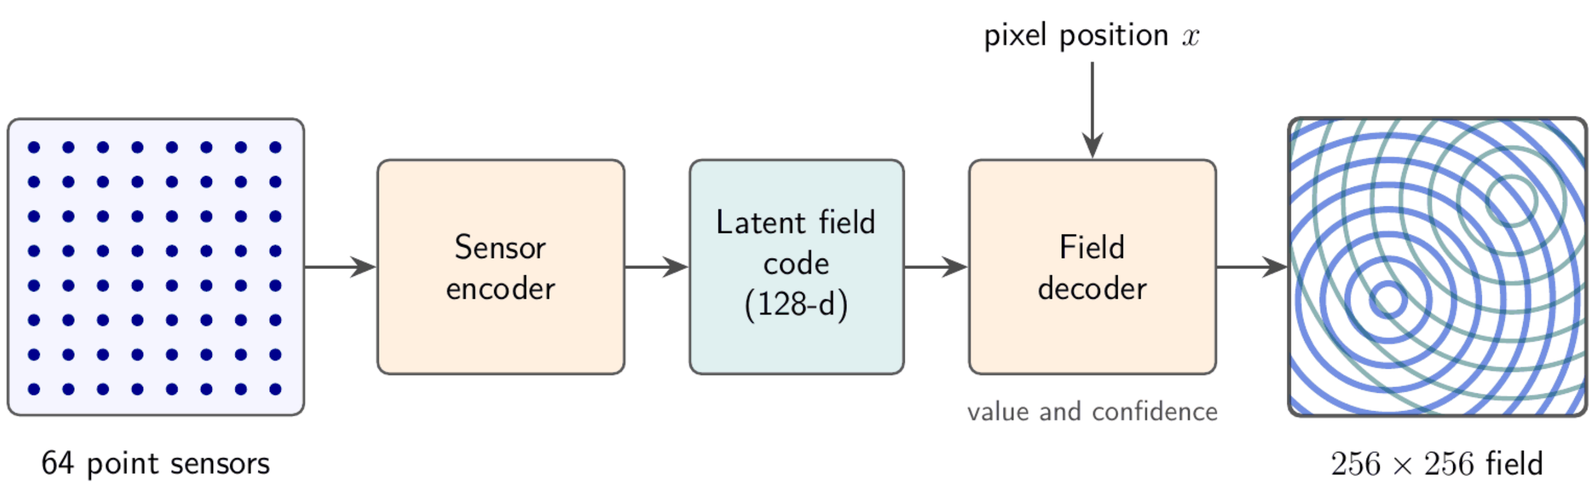
Sparse-Sensor Reconstruction of Synthetic Ripple Fields

Ada Placeholder, Ben Example, Cleo Sample

Example University (fictional), Department of Imaginary Sciences

FCSV 2026 (Fictional Conference on Synthetic Vision)

Overview



Problem

Recover a full 256×256 ripple field from 64 point sensors. Smooth interpolation loses the wave fronts between sensors.

Method

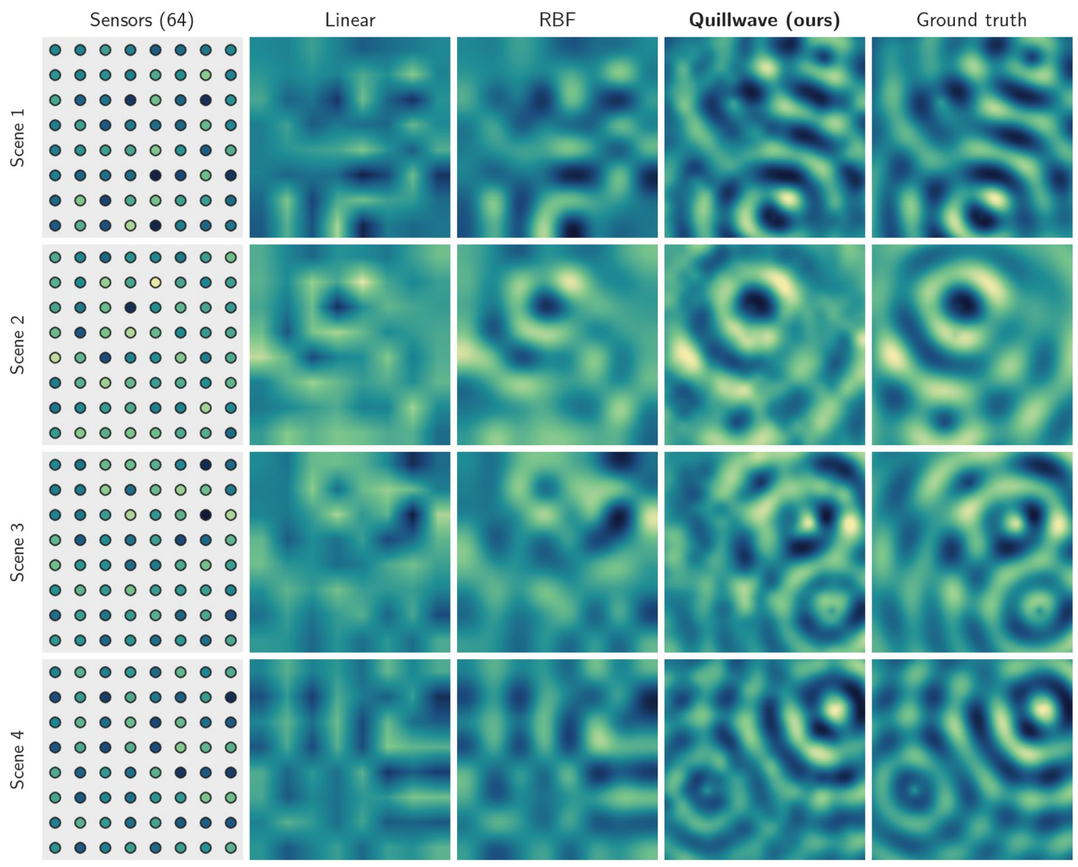
A sensor encoder maps the readings to a 128-dimensional latent field code; a coordinate-based field decoder predicts the value and a confidence at every pixel.

Main result

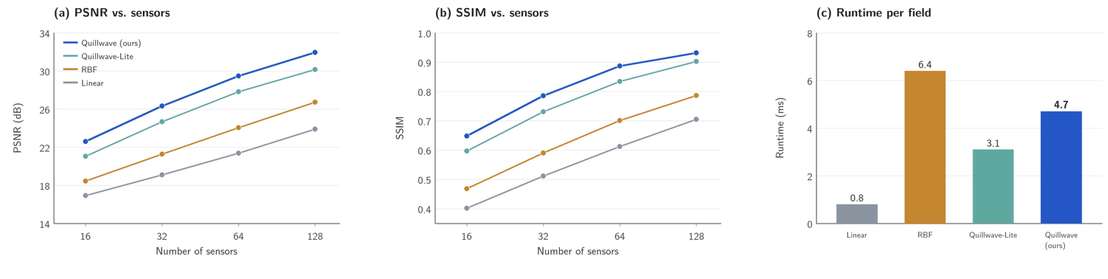
Method	PSNR (dB)	SSIM	Runtime (ms)
Linear interpolation	21.37	0.612	**0.8**
Radial-basis interpolation	24.05	0.701	6.4
Quillwave-Lite (ablation)	27.82	0.834	3.1
Quillwave (ours)	**29.46**	**0.887**	4.7

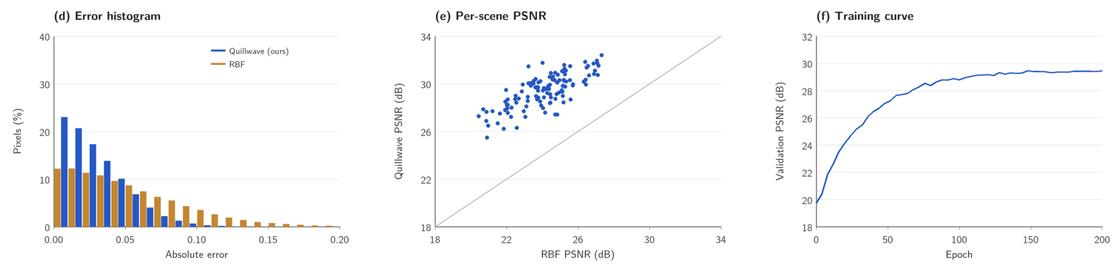
+5.41 dB over radial-basis interpolation (71.2 % lower MSE), 1.36× faster.

Qualitative comparison



Results





Limitations

- **Baselines are not compute-matched.** Quillwave is trained for 3.5 GPU-hours, while the interpolation baselines need no training at all.
- **Synthetic fields only.** Every result is on RippleBench-Synthetic; no real sensor data was used, and behaviour on real measurements is unknown.
- **Confidence is not calibration.** The per-pixel confidence map that Quillwave outputs has not been checked against observed error rates, so it should not be read as a calibrated probability.

Links

QR code placeholder
(no target)

Project page:
`https://quillwave.example.org/`

Synthetic test material. All people, institutions, venues and results are fictional.