Final Year Project 2025–2026

# Quillwave: Sparse-Sensor Reconstruction of Synthetic Ripple Fields

Ada Placeholder  Ben Example  Cleo Sample

Supervisor: Dr. Dana Specimen

Example University (fictional), Department of Imaginary Sciences

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

# Acknowledgements

We thank our supervisor, Dr. Dana Specimen, for weekly guidance, and the (fictional) Department of Imaginary Sciences for computing time. We also thank the 20 participants of our user study.

# Abstract

Quillwave reconstructs a full $256 \times 256$ ripple field from 64 point sensors. Ripple fields are procedurally generated sums of interfering circular waves: cheap to synthesise, but hard to interpolate, because the wave fronts between sensors are lost by smooth interpolation. Quillwave encodes the sensor readings into a 128-dimensional latent field code and decodes it with a coordinate-based field decoder. On RippleBench-Synthetic, a benchmark of 500 synthetic test fields generated for this work, Quillwave reaches 29.46 dB PSNR and 0.887 SSIM, 5.41 dB above radial-basis interpolation (the best baseline), which is a 71.2 % lower mean squared error. It is also 1.36× faster than radial-basis interpolation at inference. All results are on synthetic fields only.

# Contents

**1 Introduction** **4**

**2 Background** **5**

**3 System design** **6**

**4 Method** **9**

**5 Experiments** **11**
5.1 Setup . . . . . 11
5.2 Results . . . . . 11
5.3 User study . . . . . 11

**6 Discussion** **14**

**7 Limitations and future work** **15**

**8 Conclusion** **16**

# Chapter 1

# Introduction

Many sensing problems share one shape: a handful of point measurements, and a wish to know the whole field between them. Smooth interpolation is the standard answer, and it works when the field is smooth. It fails for fields made of travelling waves, where the interesting structure, the wave fronts, lies exactly between the sensors.

We study this failure on *ripple fields*: sums of damped circular waves from a few random sources, sampled on a $256 \times 256$ grid. They are easy to generate in any quantity, have a known ground truth, and carry no personal or proprietary data. We make three contributions:

- RippleBench-Synthetic, a procedurally generated benchmark with 2000 training and 500 test fields;
- Quillwave, an encoder–decoder that predicts the full field from 64 sensors and a per-pixel confidence;
- an evaluation against linear and radial-basis interpolation and an ablation, Quillwave-Lite.

# Chapter 2

# Background

Interpolating sparse probes of wave-like fields has been studied with hand-crafted bases [1, 2] and, more recently, with learned decoders [3]. Sensor placement in toy domains is discussed in [4]; we fix the layout to isolate the reconstruction step. Purely synthetic benchmarks [5] trade realism for control, and we follow the error measures recommended in [6].

# Chapter 3

# System design

The project has two parts: a training and evaluation system, and a web dashboard that shows live reconstructions from a simulated sensor array. The dashboard streams a new synthetic field every few seconds, samples it at the sensor positions, and shows the reconstruction next to the ground truth.

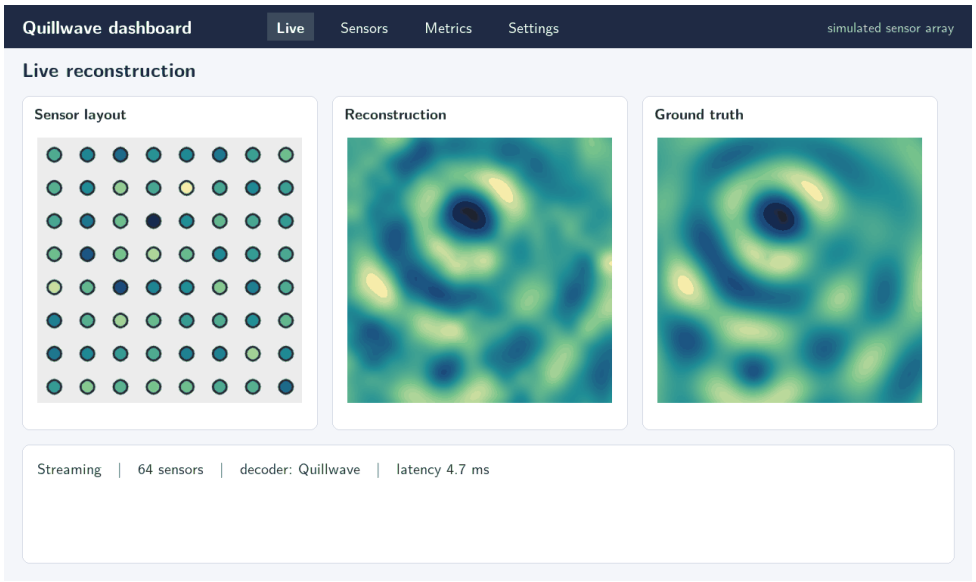



Figure 3.1: Dashboard main view: live field, sensor layout and reconstruction.

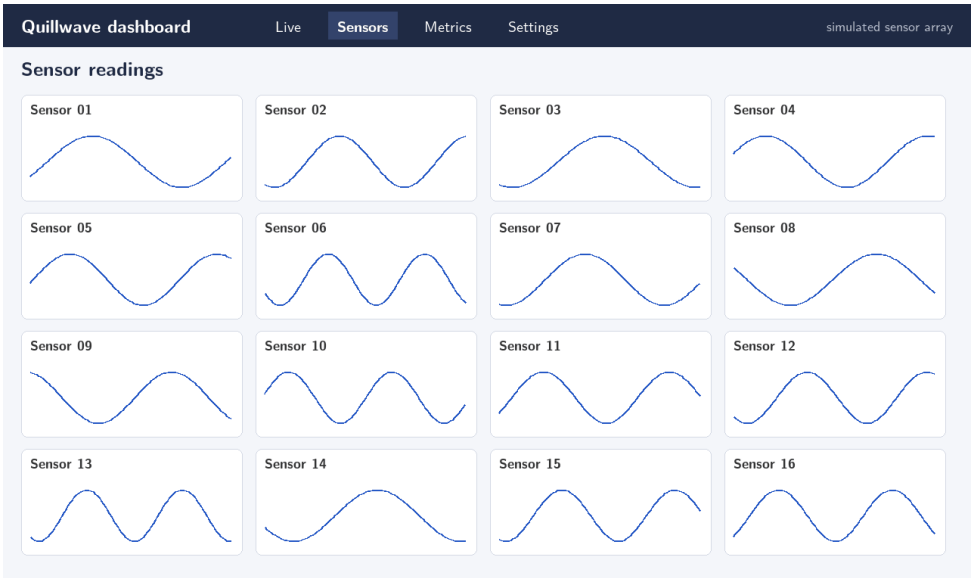



Figure 3.2: Dashboard sensor view: per-sensor readings over time.

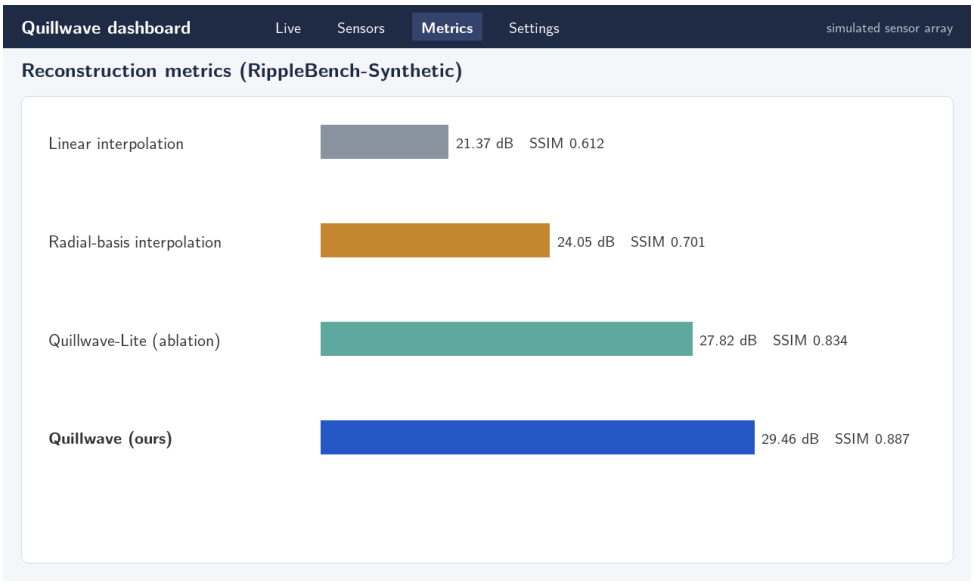



Figure 3.3: Dashboard metrics view: PSNR and SSIM of the live reconstruction.

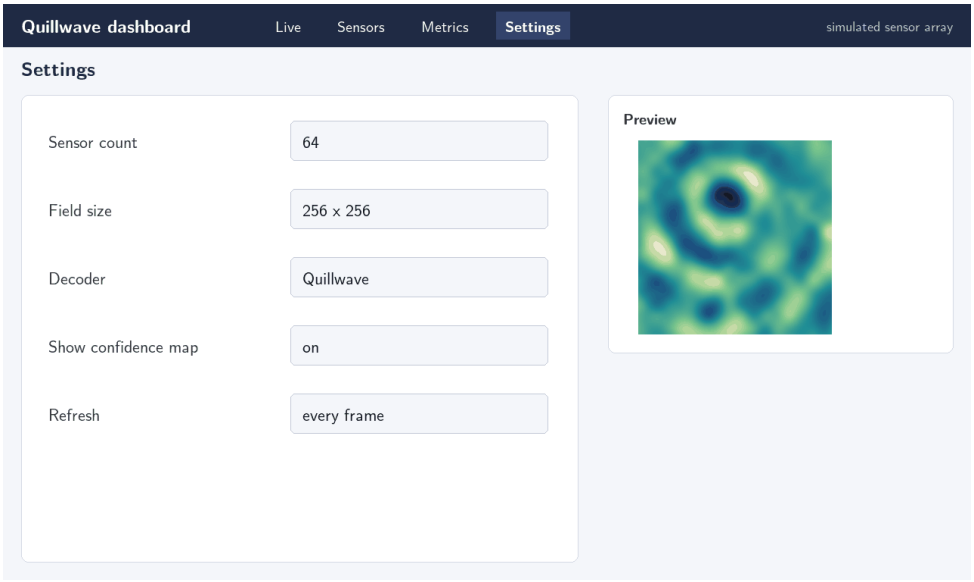



Figure 3.4: Dashboard settings view: sensor count and decoder options.

# Chapter 4

# Method

Let $s \in \mathbf{R}^{64}$ be the sensor readings at known positions $p_i$ on a regular $8 \times 8$ grid. The sensor encoder $E_\phi$ maps the pairs $(p_i, s_i)$ to a latent field code $z = E_\phi(s) \in \mathbf{R}^{128}$. The field decoder $D_\theta$ is coordinate-based: for any pixel position $x$ it predicts a value and a confidence,

$$(\hat{f}(x),\, c(x)) = D_\theta(x, z). \tag{4.1}$$

Training minimises the mean squared error between $\hat{f}$ and the ground-truth field over all pixels, plus a term that asks $c(x)$ to rank pixels by their error. Quillwave-Lite removes the latent code and conditions the decoder on the raw readings directly. The full model has 1.2M parameters and is trained for 200 epochs.

**Encoder.** Each sensor contributes a token made of its position and its reading. The tokens pass through a small set-encoder with shared weights followed by attention pooling, so the code does not depend on the order in which sensors are listed. The pooled vector is projected to the latent field code.

**Decoder.** The decoder receives a pixel position, encoded with sinusoidal features, together with the latent field code. A stack of fully connected layers with skip connections predicts two outputs: the field value and a confidence score. Because the decoder is queried per pixel, the output resolution is not tied to the training resolution, although every result in this document uses the benchmark resolution.

**Confidence.** The confidence head is trained with a ranking objective: pixels with a larger reconstruction error should receive a lower score. It is meant to show *where* the reconstruction is least reliable, not *how* reliable it is; Section 7 returns to this point.

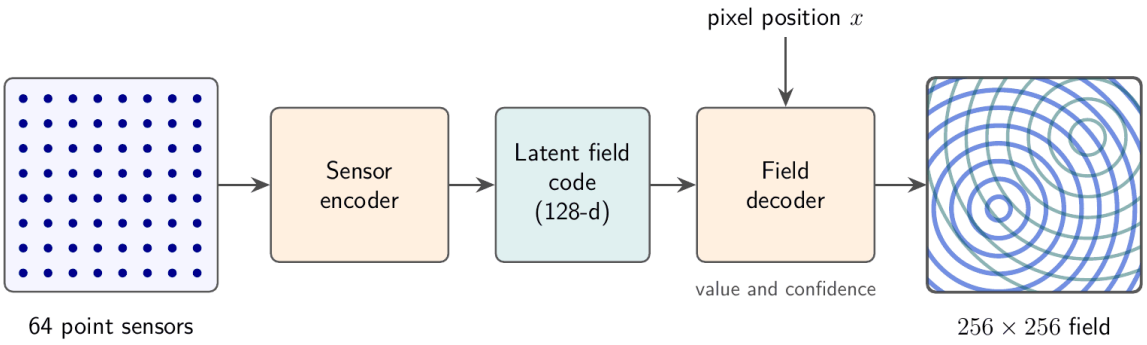



Figure 4.1: Quillwave pipeline: sensor readings are encoded into a latent field code, which a coordinate-based decoder turns into the full field.

# Chapter 5

# Experiments

## 5.1 Setup

RippleBench-Synthetic fields are $256 \times 256$ sums of damped circular waves from up to four sources with random positions, wavelengths and phases. We use 2000 training fields and 500 test fields, and 64 sensors unless stated otherwise. We compare against linear interpolation and radial-basis interpolation from the same sensors, and against the Quillwave-Lite ablation. Learned models are trained 5 times with different seeds; we report the mean over the 5 training runs. Runtime is per field at inference.

## 5.2 Results

Table 5.1 gives the main result. Quillwave reaches 29.46 dB PSNR and 0.887 SSIM, 5.41 dB above radial-basis interpolation, the best baseline; in terms of mean squared error this is a 71.2 % reduction. At 4.7 ms per field it is 1.36× faster than radial-basis interpolation (6.4 ms), though slower than linear interpolation (0.8 ms). The ablation Quillwave-Lite reaches 27.82 dB, so the latent field code accounts for part of the gain.

Table 5.2 varies the number of sensors. Quillwave is the best method at every sensor count, from 22.58 dB with 16 sensors to 31.94 dB with 128 sensors. Figure 5.1 summarises these trends together with error statistics and the training curve, and Figure 5.2 shows reconstructions of four test scenes.

**Failure cases.** Quillwave is least accurate where two wave sources lie close together between the same group of sensors: the interference pattern then varies faster than the sensor spacing, and the decoder tends to merge the two sources into one. Radial-basis interpolation fails in the same places, but more smoothly. Near the border of the field, where sensors are present only on one side, all methods lose accuracy; Quillwave keeps the wave fronts but sometimes shifts their phase.

## 5.3 User study

A small user study with 20 participants rated the side-by-side comparison videos. Mean ease of use was 4.10 and mean perceived quality 3.90 on a scale up to 5. Raw responses are not published.

Table 5.1: Reconstruction quality on RippleBench-Synthetic (64 sensors, $256 \times 256$ fields, mean over 5 training runs). Higher PSNR/SSIM is better; lower runtime is better.

| 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 |

Table 5.2: PSNR (dB) on RippleBench-Synthetic as the number of sensors changes ($256 \times 256$ fields).

| Method | 16 sensors | 32 sensors | 64 sensors | 128 sensors |
|---|---|---|---|---|
| Linear interpolation | 16.92 | 19.10 | 21.37 | 23.88 |
| Radial-basis interpolation | 18.44 | 21.26 | 24.05 | 26.71 |
| Quillwave-Lite (ablation) | 21.03 | 24.67 | 27.82 | 30.15 |
| Quillwave (ours) | **22.58** | **26.31** | **29.46** | **31.94** |

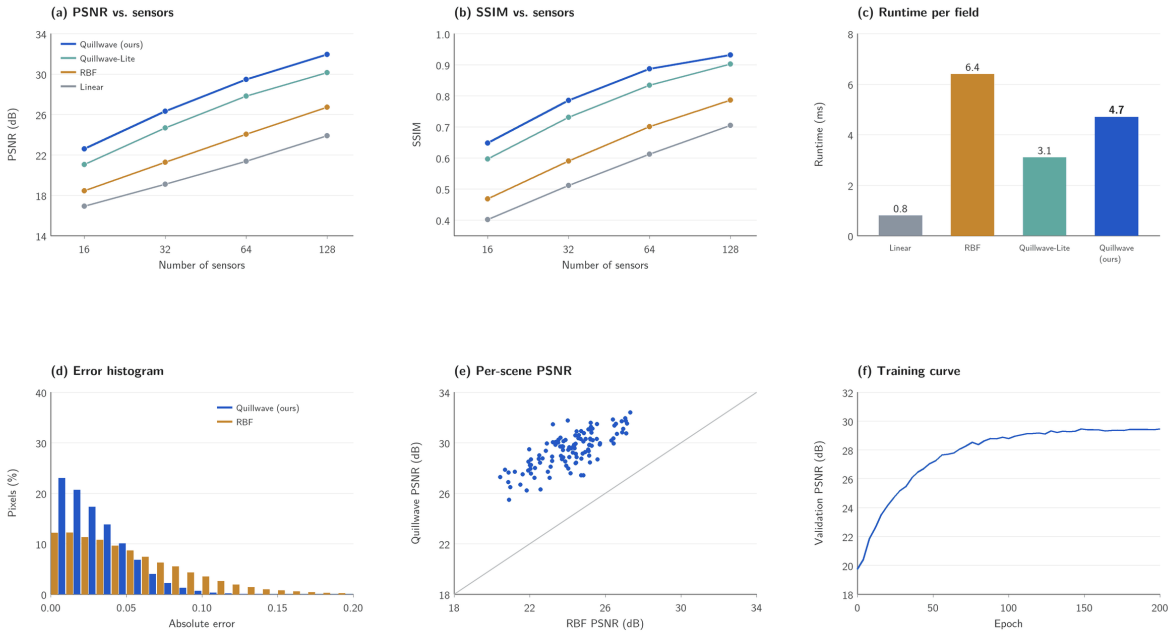



Figure 5.1: Results on RippleBench-Synthetic. (a) PSNR and (b) SSIM against the number of sensors; (c) runtime; (d) absolute error histogram; (e) per-scene PSNR, Quillwave against radial-basis interpolation; (f) validation PSNR during training.

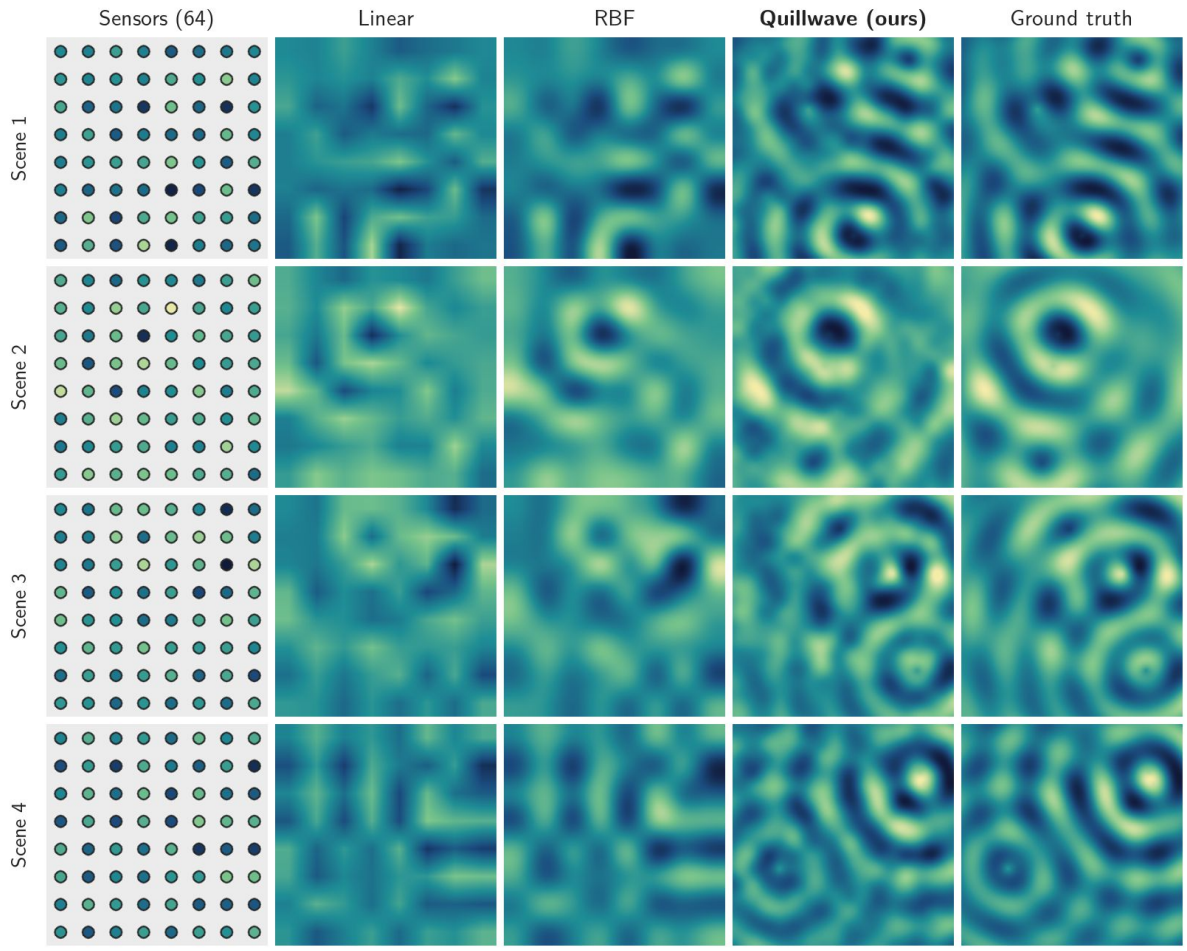



Figure 5.2: Four test scenes: sensor map, linear interpolation, radial-basis interpolation, Quillwave, and ground truth.

# Chapter 6

# Discussion

The results suggest that the main benefit of a learned latent code is a prior over plausible wave patterns: with few sensors, the decoder fills the gaps with wave fronts that are consistent with the readings, where interpolation fills them with a smooth surface. This prior is also the main risk. Because RippleBench-Synthetic is generated by the same family of processes used for training, the benchmark rewards exactly this prior, and it says little about fields from other processes. A fair comparison on real measurements would need baselines with access to the same training data and the same compute budget, and a separate check of whether the confidence scores match observed errors.

# Chapter 7

# Limitations and future work

- **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.

Future work includes testing on real sensor data and calibrating the confidence map.

# Chapter 8

# Conclusion

Quillwave shows that a learned latent field code recovers wave fronts that smooth interpolation misses, at a lower runtime than radial-basis interpolation. Whether this carries over to real measurements is open; the limitations above apply to every number in this document.

# Bibliography


[1] Q. Nobody and R. Anyone. Interpolating ripples from very few probes. *Proc. FCSV 2024*, 2024. (Fictional reference.)

[2] R. Anyone. Radial bases for sparse synthetic fields. *J. Imaginary Sci.*, 2023. (Fictional reference.)

[3] S. Someone and Q. Nobody. Learned decoders for synthetic wave fields. *Proc. FCSV 2025*, 2025. (Fictional reference.)

[4] T. Whoever. A note on sensor placement in toy domains. *J. Imaginary Sci.*, 2024. (Fictional reference.)

[5] U. Nameless and S. Someone. Benchmarks without data. *Proc. FCSV 2025*, 2025. (Fictional reference.)

[6] V. Anybody. On measuring reconstruction error. *J. Imaginary Sci.*, 2025. (Fictional reference.)