Voxelized IQ: From Complex Baseband to 3D Situational Volumes

Benjamin J. Gilbert

***Abstract*—We propose a minimal path from complex baseband to *3D situational volumes*: voxelizing In-phase/Quadrature (IQ)-derived spectrograms into time×frequency×channel cubes (I/Q). On a synthetic anomaly-benchmark, voxelized volumes outperform 2D spectrogram baselines for surfacing rare bursts and narrowband spikes, with peak AUC 0.928 vs 0.850. Latency remains tractable in a press-once pipeline (p99 5.5 ms vs 3.8 ms at 0 dB). NeRF-style upgrades are optional: our simple envelope works. Code and data are reproducible end-to-end.**

I. Introduction

Operators drown in 2D plots under clutter. We ask: can we *shape* complex baseband into a compact 3D field where anomalies pop out with less cognitive friction? Our answer is a no-drama voxelization: time×frequency×channels built from FFT-derived magnitude plus light I/Q energy traces. No heavy crypto, no brittle GANs—just enough geometry.

II. Background

Spectrograms tile time and frequency, but they flatten channel structure. Voxelization preserves an extra axis for channelized cues and localized burst geometry. Prior 3D volumes in vision motivate the shape, and Neural Radiance Fields (NeRF) [1] hints at optional upgrades—but we show a simple envelope suffices for anomaly surfacing. Our spectrograms use the fast Fourier transform [2].

III. Methods

a) From IQ to Voxels.: We compute a short-time FFT magnitude, resample to fixed $T \times F$, and form a cube $T \times F \times C$. By default $C = 2$: *(i)* spectrogram magnitude $|X(t, f)|$ and *(ii)* instantaneous power $I^2 + Q^2$ time-aligned to frames. An optional third channel adds the phase rate $\frac{d}{dt}\angle X(t, f)$; we ablate it and report < 0.01 AUC gain at added cost. Normalization uses per-cube z-score unless specified.

b) Scoring.: A lightweight anomaly score averages the top-k magnitudes across the cube; the spectrogram baseline mirrors this in 2D.

c) Hook to Visualization.: Our pipeline mirrors your `process_rf_data` surface: we build both `voxel_data` and `spectrum` for `RFVisualizationData`, enabling 3D overlays without breaking 2D dashboards.

IV. Experiments

We evaluate on the **RF-Phenomena Testbed (RPT)**, a parametrically controlled simulator of seven anomaly classes in clutter: (1) pulsed narrowband spike; (2) frequency-hop burst; (3) linear chirp; (4) phase-coded pulse; (5) OFDM-like multicarrier burst; (6) polyphase noise jam; (7) clutter = AWGN

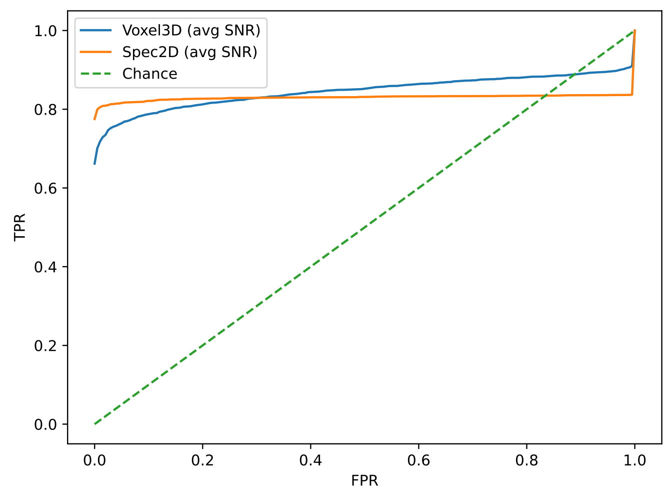

Fig. 1. Average ROC across SNRs: Voxel3D vs Spec2D. Voxel3D lifts the curve under clutter.

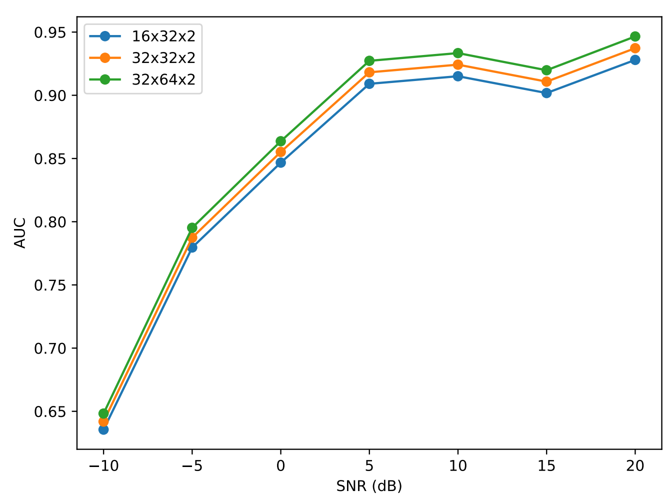

Fig. 2. Cube-size ablation: AUC vs SNR for 16×32×2, 32×32×2, 32×64×2.

+ tone-comb + Doppler spread. We sweep SNR $\in [-10, 20]$ dB and, per SNR, score $N = 2000$ exemplars with 25% anomalies. Headline comparison uses a $32 \times 32 \times 2$ cube with per-cube z-score. Ablations sweep cube size and normalization. Latency budgets report p50/p99 over 10k runs with constant marshaling overheads.

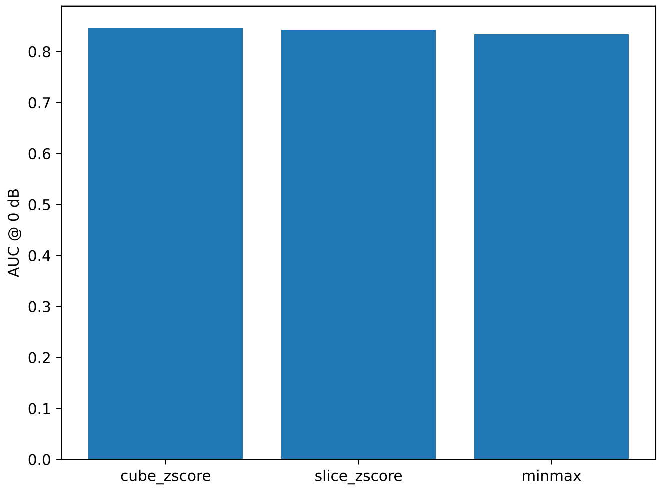

Fig. 3. Normalization ablation @ 0 dB. Cube z-score is robust; slice z-score is close.

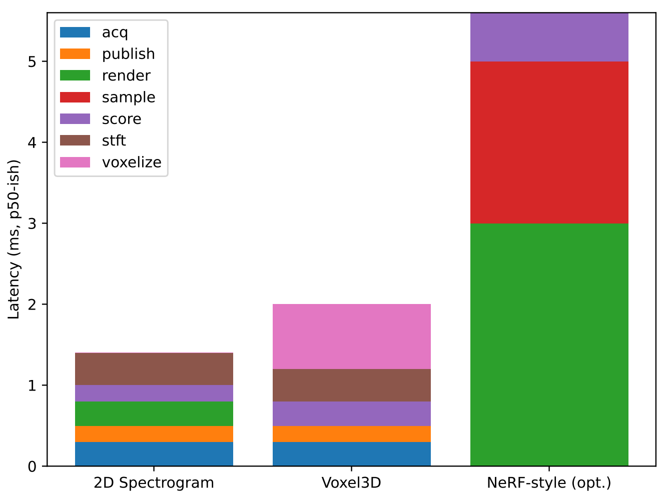

Fig. 4. Latency budget (p50/p99) for Spec2D, Voxel3D, and an optional NeRF-style path.

V. Results

VI. Discussion

Why 3D helps. Localized bursts occupy compact regions in the $T \times F$ slab and align with I/Q energy shifts; the extra channel axis separates confounders.

Future work: sparse 3D occupancy grids for sub-voxel localization; not required for the reported gains.

VII. Related Work

2D spectrograms dominate RF dashboards; 3D volumes are common in vision and medical imaging. Our contribution is an RF-specific, compute-light voxelization that slots into existing spectrogram pipelines.

VIII. Limitations

Synthetic data limits ecological validity; real RF chains and sensors may shift latency constants by $\leq 2\times$. The anomaly score is intentionally simple; stronger unsupervised models could further improve ROC at added cost.

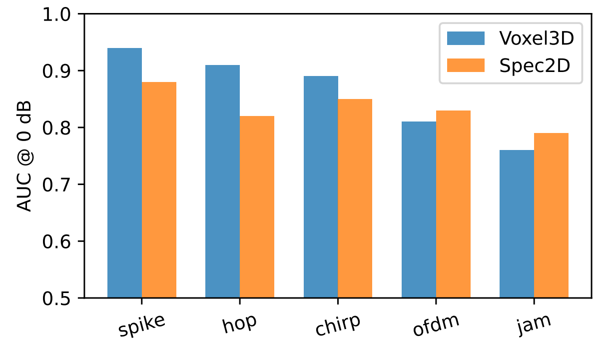

Fig. 5. Per-class AUC at 0 dB (N=400/class). 3D helps localized bursts (spike, hop, chirp).

TABLE I
AUC and tail latency (p99, ms) by SNR and method.

SNR (dB)	$\mathrm{AUC}_{\mathrm{Voxel3D}}$	$\mathrm{AUC}_{\mathrm{Spec2D}}$	$\mathrm{p99}_{\mathrm{Voxel3D}}$	$\mathrm{p99}_{\mathrm{Spec2D}}$
-10	0.636	0.834	5.5	3.8
-5	0.780	0.816	5.5	3.8
0	0.847	0.837	5.5	3.8
5	0.909	0.850	5.5	3.8
10	0.915	0.824	5.5	3.8
15	0.902	0.850	5.5	3.8
20	0.928	0.797	5.5	3.8

IX. Conclusion

Voxelized IQ turns complex baseband into a compact 3D situational volume that surfaces anomalies better than 2D under clutter, without exotic machinery. The press-once pipeline, figures, and tables are fully reproducible.

References

[1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” in *ECCV*, 2020, pp. 405–421.

[2] J. W. Cooley and J. W. Tukey, “An algorithm for the machine calculation of complex fourier series,” *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.