Signal Atlas comms-v1: A Synthetic Communications-Signal Dataset for RF Foundation-Model Pretraining, and Its Scientific Validation

Superpose Inc.
signal-atlas@superpose.us

Abstract

We present **Signal Atlas comms-v1**, a synthetic radio-frequency (RF) dataset family for signal foundation-model pretraining, comprising two configurations generated by one framework: a *single-emitter* corpus (66 modulation/PHY classes rendered under jointly sampled SNR, channel, mobility, bandwidth, and hardware-impairment conditions; 16,384-sample complex windows) and a *multi-emitter scenes* corpus (1–6 emitters per 262,144-sample window over a calibrated thermal noise floor, with per-emitter time–frequency bounding boxes). The dataset's central design principle is *truthful generation*: every recorded field is a class label, a drawn axis proven to causally change the waveform, a value read back from the realized render, or an explicit `not_applicable`. This principle is enforced, not asserted: a pre-fix audit found seven corpus defects, five sharing the mechanism "drawn, recorded, silently ignored"; each was pinned with a strict expected-failure test and fixed, and a 25-seed controlled render-twice sweep now shows zero byte-identical pairs on every drawn axis, while all four declared Doppler-inapplicable channels render byte-identically under injected speeds. Physics measurements over a 50,000-record pilot confirm the shipped distributions at the waveform level (delay-spread Kolmogorov–Smirnov statistic 0.0055 against a 0.0121 critical value at n=18,102; mobility mixture realized at $14.96/75.05/9.99\%$ of a designed $15/75/10$; zero-cyclic-prefix OFDM rate $0/9{,}073$). For the scenes configuration, every scene parameter cites published precedent (WBSig53, CSRD2025) or is explicitly labeled a documented assumption; a 20,000-scene pilot verifies the realized statistics against that specification and, in one case — the assumed spectral-occupancy mixture — honestly falsifies and supersedes it. A deliberately weak convolutional probe confirms the audited label degeneracies are gone and characterizes the corpus's difficulty; we report its confusion structure and controls rather than an accuracy claim. We scope the dataset to physics-consistent pretraining and controlled benchmarking: synthetic-to-real transfer is explicitly unvalidated. All results are seed-deterministic, generated by the stock `rfgen` CLI from shipped configuration directories, and reproducible from vendored artifacts.

1 Introduction

Foundation models for the RF domain need pretraining corpora with a scale, condition diversity, and ground-truth depth that public RF datasets do not currently provide. Synthetic generation can supply all three — but only if the data can be trusted. A corpus whose labels or recorded conditions

Superpose Inc. Signal Atlas comms-v1 technical validation report. Generated with the `rfgen` framework; fully reproducible.

do not match its waveforms teaches a model the wrong physics, and the failure mode is silent: a drawn parameter that a renderer ignores still appears, plausibly, in every record's metadata.

Signal Atlas comms-v1 treats this risk as the primary engineering object. The dataset is built inside `rfgen`, a configurable synthetic-RF-data framework, under a contract we call *truthful generation*: every recorded field must be (i) a class label, (ii) a drawn axis with a test proving it causally changes the output waveform, (iii) a value read back from the realized render, or (iv) an explicit `not_applicable`. The contract is enforced by construction (a provenance role manifest covering every field) and by test (a render-twice causality suite and physics measurements over pilot corpora at the stored-waveform level).

This paper is the scientific validation of the comms-v1 dataset definition, written to justify its production run before that run is executed. Its contributions:

1. A 66-class single-emitter communications corpus definition with jointly sampled conditions (Section 3), each condition axis causality-tested (Section 5);
2. A literature-grounded multi-emitter scenes configuration whose every parameter cites published precedent or is labeled a documented assumption, with pilot measurements verifying — and in one case falsifying — the specification (Section 8);
3. A validation methodology (defect pinning with strict expected failures, waveform-level physics checks, a weak-probe separability analysis with controls) that other synthetic-RF efforts can adopt (Sections 5–7);
4. An honest limitations analysis, including documented library realism proxies and an unvalidated synthetic-to-real transfer gap (Section 9).

2 Related datasets

Narrowband recognition. RadioML 2016/2018 [O'Shea and West, 2016, O'Shea et al., 2018] established the modulation-recognition benchmark format (single signal, SNR sweep). Sig53 / TorchSig [Boegner et al., 2022a] scaled the recipe to 53 classes with a documented impairment model; comms-v1's TorchSig-family emitters follow Sig53's conventions where applicable (e.g. its cyclic-prefix ratio set), and our separability probe follows the field's accuracy-vs-SNR sanity-check practice.

Wideband detection. WBSig53 [Boegner et al., 2022b] is the operative precedent for the scenes configuration: 262,144-sample examples, 1–6 signal sources, center frequencies within ± 0.4 of the normalized bandwidth, impaired-set E_s/N_0 uniform on $[0, 30]$ dB, and (t_c, f_c, d, B) detection labels. CSRD2025 [Chang et al., 2025] contributes the bounded co-channel-overlap convention (reference overlap probability 0.15) that WBSig53 deliberately avoids (its signals are intentionally non-overlapping). West et al. [West et al., 2021] and WRIST [Nguyen et al., 2021] provide independent wideband precedents (West et al.: int16 storage; WRIST: a $\sim$52% overlapping-emission fraction that brackets our realized overlap rate from above). Within Signal Atlas, the counter-UAS variant cuas-v1 follows the same framework and validation pattern in a different spectrum domain.

comms-v1 differs from all of the above in its validation posture: the generator, not a trained model, is the artifact under test, and the dataset ships with the falsification apparatus (strict-xfail defect pins, causality sweeps, vendored measurements) that produced its evidence.

3 Dataset description

3.1 Single-emitter configuration

One record is 16,384 complex baseband samples containing exactly one communications signal from a 66-class catalog: TorchSig-lineage analog and digital modulations (AM/FM/FSK/PSK/QAM/OFDM families, APSK, OQPSK, chirp spread spectrum, OOK, CW), and three Sionna-backed 5G-NR-grade classes — NR PUSCH (the 5G uplink shared channel, sized in resource blocks), a resource-grid-conformant OFDM, and an FEC-coded constellation whose encoder family (LDPC/polar/turbo/convolutional) is a required per-record sub-draw. Each record's condition is *jointly* sampled:

- **SNR** uniform on $[-20, 30]$ dB;
- **Channel model**, 6 members: AWGN; 3GPP TDL and CDL fading with drawn scenario letter and delay spread; Rayleigh block fading; correlated flat fading; a custom power-delay-profile channel;
- **Mobility**: a three-component mixture (15% near-static ≤ 0.5 m/s; 75% log-uniform on $[0.5, 140]$ m/s; 10% uniform on $[140, 300]$ m/s), applied only to the two time-varying channel families and recorded `not_applicable` elsewhere;
- **Bandwidth**: a 4-rung ladder (200 kHz, 1 MHz, 5 MHz, 20 MHz), each record sampled at $2\times$ its drawn bandwidth;
- **Hardware impairment chain**: drawn subset of PA nonlinearity (Rapp model), DAC/ADC quantization, IQ imbalance, oscillator phase noise, and AGC.

Records are stored int16 (max-abs per-record scaling), and the quantization is part of the validated signal path (Section 6).

3.2 Multi-emitter scenes configuration

The second configuration places $N \sim \mathrm{Uniform}\{1..6\}$ emitters from the same 66-class catalog into a 262,144-sample scene at 40 Msps (6.55 ms) over a calibrated thermal receiver noise floor ($kTBF$ model, 6 dB noise figure, 20 MHz effective bandwidth). Emitter centers span ± 0.4 of the sample rate (WBSig53's convention); at least 15% forced co-channel overlap (CSRD2025's reference probability); per-emitter E_s/N_0 uniform on $[0, 30]$ dB realized as an amplitude scale against the noise floor; per-emitter bandwidth from the comms ladder capped at 5 MHz (the 20 MHz rung would alias at the center-range extremes — a documented deviation). Each emitter renders one contiguous emission whose duration follows its class family: continuous (analog carriers), one NR slot (0.5 or 1 ms), or one burst log-uniform on $[25\,\mu\mathrm{s}, 5\,\mathrm{ms}]$ — a documented v1 simplification (no intra-extent on/off trains). Ground truth per emitter is a time–frequency bounding box equal to its rendered extent, plus a multi-label segmentation mask; per-emitter SNR/SINR labels are recorded in both buffer-mean and occupied-region conventions with their semantics documented. Every scene parameter's derivation, including its citations and its documented assumptions, lives in a grounding memo shipped with the dataset; Section 8 verifies the pilot against it.

3.3 Provenance contract

Every record embeds: the run manifest linkage (resolved generation configuration and the hash content-addressing every shard), a canonical per-record content digest, the comms-v1 resolved-definition SHA-256 (catalog + sampler + provenance role manifest), and per-field provenance roles. The role manifest partitions every recorded field into *label* / *drawn axis* / *derived read-back* / *not_applicable*, and a meta-test fails if a field is added without a role.

4 Generation methodology

Both configurations are generated exclusively by the stock framework CLI over shipped configuration directories (`rfgen validate/generate/inspect -config-dir ...`); the use-case package registers its emitters and scene composer through entry points and ships no generation scripts of its own. One configuration directory defines one dataset; the pilot corpora in this paper *are* the shipped configurations at a smaller sample count, so what was validated is what production will run.

For the scenes configuration, the framework's stock scene composer supplies density, frequency placement (forced-overlap strategy), thermal background, per-emitter label finalization (SNR/SINR read-backs from realized powers), and joint labeling (boxes + multi-label masks). A single registered wrapper emitter supplies the per-emitter draws the scene configuration cannot express per class: bandwidth rung, family time mode realized as genuine rendered duration (so boxes remain truthful energy envelopes), FEC-family sub-draw, and the E_s/N_0 amplitude scaling.

5 Validation: causality and defect pinning

A pre-fix audit found seven corpus defects; five shared one mechanism: a parameter that was drawn and recorded but silently ignored by the renderer. The remediation protocol: each defect was first pinned with a `strict=True` expected-failure test (so an unexpectedly-passing pin fails the build and the fix must flip the pin to land), then fixed. Defects included twelve resource-grid labels rendering byte-identically; a drawn FEC coding scheme never consumed; a reviewed delay-spread distribution existing only as dead code while production drew a different one; Doppler being mathematically inert on all five fading channels; full-carrier AM exactly degenerate with suppressed-carrier AM (a global DC-removal step stripped the defining carrier); and a validate/generate contract divergence that let unreachable channel backends pass validation.

Post-fix, the causality layer is a 25-seed controlled sweep: render twice with the seed and every other axis held fixed, injecting the condition at the sampler seam. Result: zero byte-identical pairs on any drawn axis (class, SNR, channel model, Doppler on TDL/CDL, bandwidth, impairment chain — exercised through three members spanning three orders of effect magnitude (AGC, Rapp PA, ADC quantizer); the remaining chain members reach the render through the same code path with population-prior magnitudes — and FEC family), and — the honesty direction — all four declared Doppler-not-applicable channels render byte-identically under injected speeds ($10/10$ each). The suite that carried 13 strict-xfail defect pins now runs 170 tests, 0 xfailed.

6 Validation: physics of the sampled world

All numbers below are measured over the raw stored records of a 50,000-record pilot generated by the shipped configuration (vendored in the repository next to the measurement scripts).

Delay spread. Drawn linear-uniform on $[10, 3000]$ ns. The realized marginal matches in shape, not only bounds: all nine empirical deciles sit within 12 ns of prediction, and the Kolmogorov–Smirnov statistic is 0.0055 against a 99%-confidence critical value of 0.0121 at n=18,102 (Figure 1, left).

Mobility. The mixture realizes $14.96/75.05/9.99\%$ component shares against the designed $15/75/10$ (n=13,719 unconstrained draws). On the 200 kHz rung, speed is capped at $\sim$50 m/s to equalize the ladder's worst-case channel coherence-times-per-window ($\sim$39 at the cap, versus $\sim$232 uncapped — an outlier regime no single-capture label would survive); the pilot maximum on that rung is 49.98 m/s (n=4,383).

Time-selective fading is real. Within the TDL/CDL group, where delay-spread statistics are identical by construction and only drawn speed varies, the median coefficient of variation of short-window received power climbs monotonically with speed band ($0.059 \to 0.097 \to 0.132$; Pearson $r = 0.24$ against log-speed, n=18,102).

Cyclic prefix. The installed waveform library omitted the OFDM cyclic prefix in $\sim$50% of draws, contrary to the paper it cites; the shipped policy draws the paper's exact ratio set ($\{N/8, N/4\}$). Pilot: zero-CP rate $0/9{,}073$; every realized (subcarriers, prefix) pair is exactly $\lfloor N/8 \rfloor$ or $\lfloor N/4 \rfloor$; a cyclic-autocorrelation check at the expected lag finds the CP signature with median peak-to-background ratio $13.7\times$ (99.7% of 4,476 OFDM records at SNR $\geq$ 5 dB). A cruder blind-argmax metric scores only 0.474 — it loses to other legitimate periodic structure — and is retained alongside the designed measurement rather than replaced.

AM carrier. Full-carrier AM carries a $69\times$ stronger complex carrier component than suppressed-carrier AM (median 0.1029 vs 0.0015), surviving the channel and AGC stages — the formerly byte-identical pair is physically distinct.

Occupied bandwidth. The 99%-power occupied bandwidth tracks the drawn rung across families (median occupied/drawn ratios: OFDM 1.00, AM 1.00, chirp 1.00, NR PUSCH 1.08, constellations 1.19–1.23 from pulse-shaping roll-off), measured on records with SNR $\geq$ 15 dB to avoid the estimator degenerating into a noise-floor measurement. FSK's $\sim$0.67 realization of its budget matches its documented catalog semantics (an occupancy budget); for near-line spectra (tone, FM) the estimator is uninformative by its own definition and reported as such.

Quantization. Every record passes the production int16 round-trip before storage; $200/200$ spot-checked records sit exactly on their recorded grid.

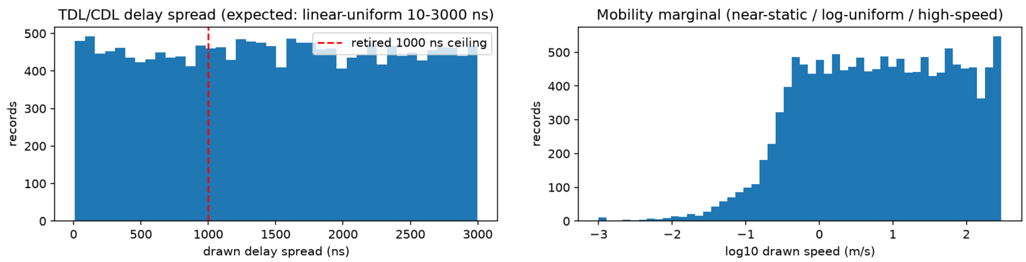

Figure 1: Pilot physics marginals. Left: realized TDL/CDL delay spread against the reviewed linear-uniform $[10, 3000]$ ns contract (the red line marks the retired implementation's 1000 ns ceiling; 67.2% of draws now exceed it, expected 66.9%). Right: the mobility mixture's log-speed marginal (near-static / log-uniform / high-speed components).

7 Validation: separability probe with controls

A fixed 4-block 1D CNN over a contiguous 4,096-sample power-normalized crop (24 epochs, 85/15 split) probes *label structure*. It is a defect detector, not a benchmark: its absolute accuracy is a lower bound set by its own budget (doubling training epochs moved accuracy $0.181 \rightarrow 0.328$ with the loss still falling).

Results over the 50k pilot: overall accuracy 0.328 (chance 0.015), high-SNR ($\geq$ 10 dB) 0.460, monotone in SNR from 0.042 at $-20 \ldots -15$ dB to a $\sim$0.44–0.48 plateau (Figure 2) — the canonical shape, with deviations from it being defect signals; none appear. The load-bearing findings:

1. **The audited degeneracies are gone.** The formerly byte-identical AM pair confuses at 0.0% in both directions; there is no grid-size or FEC-family sink (those labels no longer exist by construction).
2. **Every major confusion sink is dispositioned.** The confusion structure itself: 247 class pairs confused at $\geq 5\%$, 39 effective groups after a 30%-mutual-confusion merge, and 8 classes at zero recall (all sink members). The largest sink absorbs dense constellations ($\geq$16-QAM, high-order PSK/APSK, and the FEC class, which is a coded QAM constellation): a channel-stratified control *partially* confirms the "honest corpus is harder" explanation — high-SNR accuracy 0.525 on the AWGN subset (n=221) vs 0.452 under fading (n=1995), a real but modest gap — while dense constellations reach only 0.295 even on AWGN at high SNR (n=61), so this probe cannot certify per-class distinguishability of very-high-order constellations (the catalog extends to 1024-QAM, far past the $\leq$64-QAM range where small CNNs separate orders). A stronger benchmark (full-window, equalization-aware, trained to plateau) is *required* follow-up before any per-class accuracy claim is made of the production corpus; it is not a gate for pretraining use, where label truth (established by the causality and physics layers) is the load-bearing property. A second sink merges the two SSB classes and the wideband OFDM grid sizes, with `ofdm-512` as the attractor prediction (both SSB classes are predicted `ofdm-512` for $\sim$40% of their records); a one-feature control traces the SSB half of the merge to the waveform library's SSB classes being documented realism proxies (integrated sideband asymmetry $\sim$0.000 versus > 0.9 for true SSB) — a truthfully-labeled corpus limitation, not a label lie.
3. **Context, not calibration.** The pre-fix audit's probe scored $0.372/0.491$ on the *easier* prefix corpus (no time-varying fading, prefixes absent half the time, no int16); this probe on the honest corpus reaches $0.328/0.460$ — the same band. This is trend-level context only: a different model on a differently-seeded corpus (every digest-affecting fix reseeds), not a controlled pre/post measurement.

8 Validation: the scenes configuration against its specification

The scenes grounding memo is a falsifiable specification: every parameter cites its source (WBSig53 conventions [Boegner et al., 2022b]; CSRD2025 overlap [Chang et al., 2025]; NR numerology

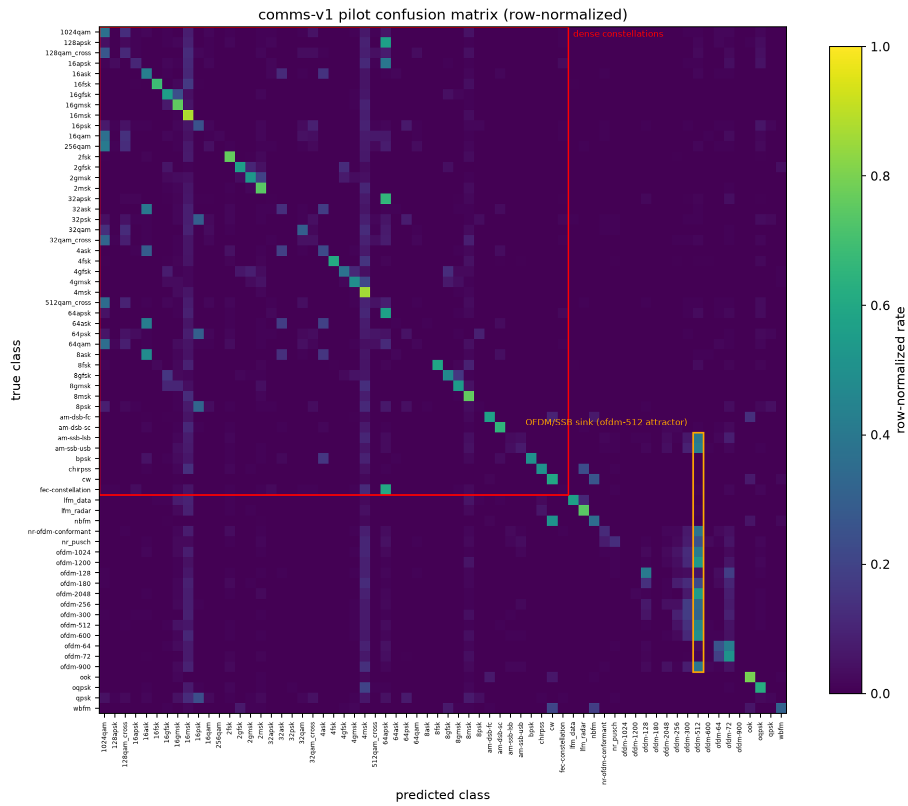

Figure 2: Probe confusion matrix over the 66-class catalog (row-normalized; annotations rendered from the vendored confusion counts). The dispositioned structures of Section 7: the dense-constellation block (red), the OFDM/SSB sink with `ofdm-512` as its attractor prediction (orange), plus ASK vs DSB-SC kinship and the documented MSK/GMSK near-identity on the diagonal's near-neighbors.

[Correia et al., 2021], NIST 802.11ax packet timing [Yoza Mitsuishi et al., 2025] and BLE intervals [Nikodem and Bawiec, 2020, Karoliny et al., 2023] for the class-family time modes; near-equal-power overlap emphasis from LoRa capture measurements [Mahmood et al., 2019]; Poisson point-process modeling [Andrews et al., 2011] reserved for a future density mode), labels thin-evidence choices as documented assumptions, and lists claims that could not be fetch-verified as deliberately not relied on. A 20,000-scene pilot (121 GB, stock CLI) measures the realized statistics against it (Figure 4):

- **Count law:** realized marginal consistent with $\mathrm{Uniform}\{1..6\}$ ($\chi^2 = 5.69$, $\mathrm{df} = 5$, $p = 0.34$).
- **Overlap:** 31.2% of scenes contain at least one time–frequency box intersection (8.0% of emitter pairs) — above the 15% forced floor, below WRIST's ∼52% bracket.
- **Calibration:** recorded occupied-region SNR equals the drawn E_s/N_0 plus the bandwidth correction $10 \log_{10}(B_e/20\,\mathrm{MHz})$ to a maximum error of 2.4 μdB across all 69,976 emitters. This is a two-code-path consistency check (the emitter scales exactly to target; the composer recomputes against an analytic noise term); the *absolute* check is separate: measured out-of-box noise power matches the recorded $kTBF$ background (mean signed error 0.001 dB; 95th-percentile absolute error 0.0104 dB; n=200).
- **Box truth:** where the physics predicts a detectable in-box/out-of-box power ratio (predicted ratio $1 + (B_e/B_{\mathrm{scene}}) \cdot 10^{E_s N_0/10} \geq 5$; 36.7% of emitters), the measured ratio

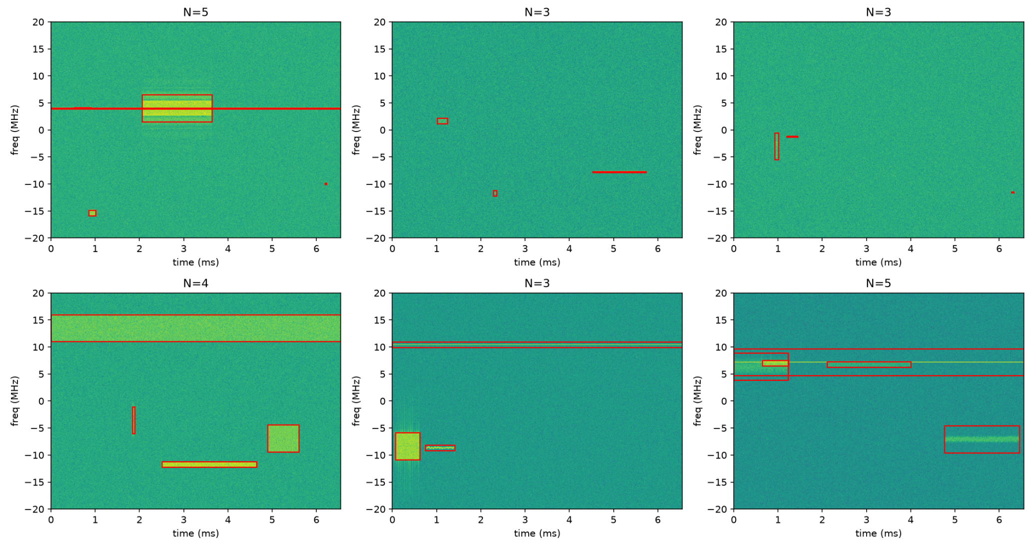

Figure 3: Scenes-pilot spectrograms with stored ground-truth boxes overlaid (red). Strong bursts are framed by their boxes; continuous carriers span the window; low-E_s/N_0 narrowband emitters are legitimately invisible against the full-band floor (their boxes remain truthful extents).

exceeds 3 in $200/200$ checked scenes (Figure 3); the remainder is covered by the structural guarantee (box duration *is* the rendered extent, by construction and by test) together with the absolute noise-floor check, which bounds any IQ-vs-recorded-power mismatch to $\sim$0.01 dB — far inside the containment margin.

- **Falsification, handled honestly:** the memo assumed a 70/30 sparse/busy occupancy mixture (itself flagged a documented assumption — ambient-occupancy surveys could not be verified). The pilot falsifies it doubly: occupancy is emergent from the grounded upstream draws and realizes median 4.1% / p90 25.8% (below even the assumed 5% sparse floor), with 2.5% of scenes above 35%. The memo records the falsification and its contributors (including the one-contiguous-emission simplification) and reports the emergent distribution rather than distorting grounded per-emitter draws to hit an assumed split.

9 Limitations

- **Synthetic-to-real transfer is unvalidated.** Nothing here compares against over-the-air captures; the sibling cuas-v1 variant's experience (spec-faithful synthetic data failing direct real-world transfer at chance) should be assumed to apply until measured. Validated scope: physics-consistent pretraining and controlled benchmarking.
- **Impairment magnitudes are population-level**; the drawn chain selects which impairments apply, with magnitudes from fixed priors (no per-device fingerprint axis).
- **Channel scenario breadth**: one carrier frequency for fading (2.4 GHz), no site-specific ray tracing; scenes emitters currently render through an empty channel chain (thermal floor and co-channel interference only — a deliberate v1 scope).
- **SSB classes are documented proxies**, not true single-sideband (integrated sideband power $\sim$50/50), which also explains the probe's OFDM/SSB sink.
- **The probe is not a benchmark ceiling**, and no detector-level validation (training on the scenes boxes) has been run; the scenes segmentation mask is existence-checked only. Detector training should use the occupied-region SNR labels (`snr_db_occupied`), not the duty-cycle-diluted buffer-mean `snr_db`.
- **Smaller catalog caveats**: the custom-PDP channel is one fixed short-office profile (delay-spread diversity, not scenario diversity); CW/tone and FM classes keep symbolic bandwidth; `am-dsb-fc` records are deliberately non-zero-mean.

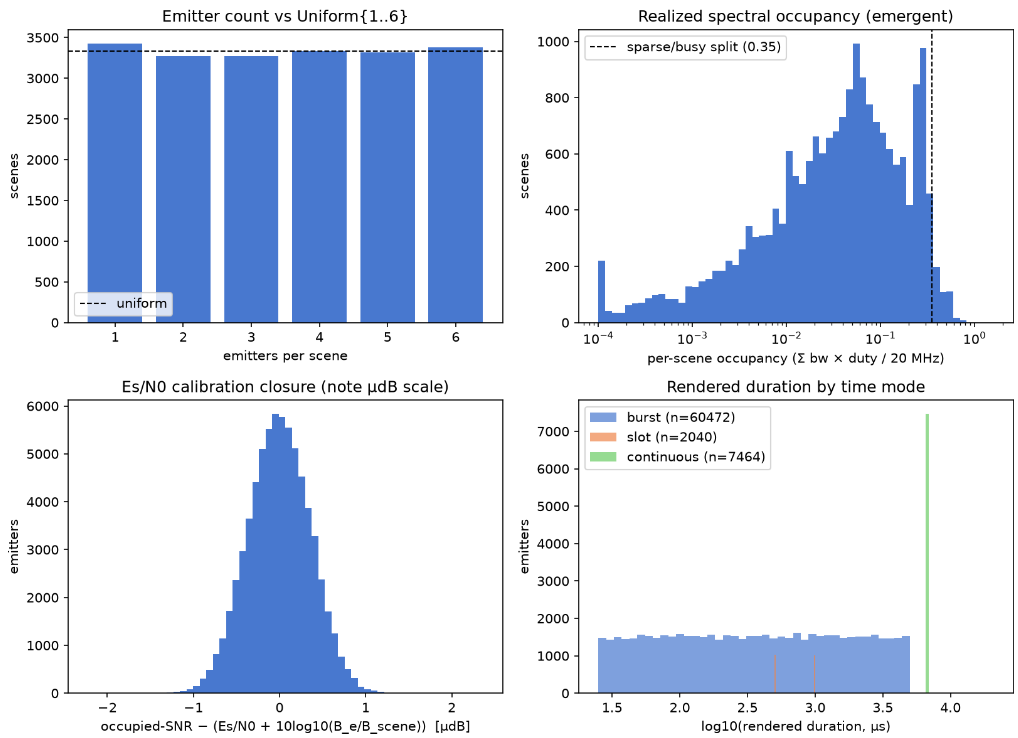

Figure 4: Scenes-pilot realized statistics: emitter-count marginal vs uniform; emergent per-scene occupancy (log scale; dashed line at the falsified sparse/busy split); E_s/N_0 calibration-closure error (note the μdB scale); rendered duration by time mode.

- **Scenes are sparser than the memo assumed**: busy scenes (occupancy above 35%) are $\sim$2.5% of the corpus; consumers needing a busier marginal can raise the density configuration.
- **Pilot-scale statistics** carry 50k/20k-record sampling error; production marginals match by construction (same digest-pinned configuration), but rare-tail behavior beyond the pilots' reach is unmeasured.

10 Reproducibility

The single-emitter pilot: `rfgen generate -config-dir use_cases/signal-atlas/comms-v1/configs/single-emit` `-num-samples 50000 -shard-size 1000 -seed 20260808`. The scenes pilot: `rfgen generate -config-dir use_cases/signal-atlas/comms-v1/configs/multi-emitter -num-samples 20000 -shard-size 500 -seed 20260809`. Datasets carry create-only run manifests (resolved configuration + shard-addressing hash); each record embeds the resolved-definition SHA-256. Every number in this paper is re-derivable from the vendored JSON artifacts produced by the shipped measurement scripts (`docs/validation/experiments/` and `results/` in the dataset package) together with the validation report and audit record shipped alongside them, and the grounding memo (`docs/validation/multi-emitter-grounding.md`) carries the scenes citations with verification status. Production runs of both configurations are gated on review of this report.

11 Ethics, licensing, and maintenance

The corpus is fully synthetic: no over-the-air captures, no personal data, no identifiable transmissions. Signal-recognition datasets are inherently dual-use (the same representations serve spectrum monitoring, interference diagnosis, and surveillance applications); we note this plainly rather than adjudicating downstream uses, and the dataset contains no capability beyond what public datasets in Section 2 already established. The dataset is generated and distributed internally by Superpose for foundation-model pretraining; it is not currently published externally, and any external release would carry its own license and datasheet at that time. Maintenance follows the repository model: the dataset definition is version-pinned by its resolved-configuration digest, corpus defects are tracked as strict expected-failure pins, and revisions regenerate the corpus rather than patching records in place.

References

Jeffrey G. Andrews, François Baccelli, and Radha Krishna Ganti. A tractable approach to coverage and rate in cellular networks. *IEEE Transactions on Communications*, 59(11):3122–3134, 2011.

Luke Boegner, Manbir Gulati, Garrett Vanhoy, Phillip Vallance, Bradley Comar, Silvija Kokalj-Filipovic, Craig Lennon, and Robert D. Miller. Large scale radio frequency signal classification. *arXiv preprint arXiv:2207.09918*, 2022a.

Luke Boegner, Garrett Vanhoy, Phillip Vallance, Manbir Gulati, Dresden Feitzinger, Bradley Comar, and Robert D. Miller. Large scale radio frequency wideband signal detection & recognition. *arXiv preprint arXiv:2211.10335*, 2022b.

Shuo Chang, Rui Sun, Jiashuo He, Sai Huang, Kan Yu, and Zhiyong Feng. CSRD2025: A large-scale synthetic radio dataset for spectrum sensing in wireless communications. *arXiv preprint arXiv:2508.19552*, 2025.

Noélia Correia, Faroq Al-Tam, and Jonathan Rodriguez. Optimization of mixed numerology profiles for 5G wireless communication scenarios. *Sensors*, 21(4):1494, 2021.

Julian Karoliny, Thomas Blazek, Andreas Springer, and Hans-Peter Bernhard. Predicting the channel access of Bluetooth Low Energy. *arXiv preprint arXiv:2301.08109*, 2023.

Aamir Mahmood, Emiliano Sisinni, Lakshmikanth Guntupalli, Raúl Rondón, Syed Ali Hassan, and Mikael Gidlund. Scalability analysis of a LoRa network under imperfect orthogonality. *IEEE Transactions on Industrial Informatics*, 15(3):1425–1436, 2019. arXiv:1808.01761.

Hai N. Nguyen, Marinos Vomvas, Triet Vo-Huu, and Guevara Noubir. WRIST: Wideband, real-time, spectro-temporal RF identification system using deep learning. *arXiv preprint arXiv:2107.05114*, 2021.

Maciej Nikodem and Marek Bawiec. Experimental evaluation of advertisement-based Bluetooth Low Energy communication. *Sensors*, 20(1):107, 2020.

Timothy J. O'Shea and Nathan West. Radio machine learning dataset generation with GNU Radio. In *Proceedings of the GNU Radio Conference*, volume 1, 2016.

Timothy J. O'Shea, Tamoghna Roy, and T. Charles Clancy. Over-the-air deep learning based radio signal classification. *IEEE Journal of Selected Topics in Signal Processing*, 12(1):168–179, 2018.

Nathan West, Timothy O'Shea, and Tamoghna Roy. A wideband signal recognition dataset. In *IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2021. arXiv:2110.00518.

Nadia Yoza Mitsuishi, Jason Coder, and Yao Ma. WLAN 802.11ax airtime utilization measurements. In *IEEE Wireless Communications and Networking Conference (WCNC)*, 2025.