# Coherent Radio-Frequency Direction Finding for Small Unmanned Aircraft: A Simulation Study of an 8-bit Five-Channel Array and a 16-bit Four-Channel Array

*With an adversarial treatment of the simulation and test methodology*

Mark Ludwikowski

CGDroneHunter Counter-UAS Program

Manuscript draft — 29 October 2026 — prepared for peer review

Code, trained weights, and a one-command reproducer: git branch **lm-detector**, `tools/rf/doa_nn/improvements/`

**Abstract.** Direction finding (DF) on the radio links of small unmanned aircraft is a core sensing task for counter-UAS. We study two coherent receiver architectures through simulation: a four-channel, 16-bit phase-coherent array of the Epiq Sidekiq X4 class [20], and a five-channel, 8-bit coherent array of the KrakenSDR class built from RTL2832U front ends [14,15,21]. On a shared impaired array manifold — per-element gain and phase error, mutual coupling, position jitter, a coherent multipath ray, finite snapshots, and additive noise — we evaluate classical subspace estimators [1,3,6], a direct neural angle classifier [12], and seven candidate enhancements, each behind a planted-fault self-test. We then add the physics that actually separates the two receivers and that the baseline manifold omitted: analog-to-digital quantization and dynamic range, calibration residual, and LO phase noise. Four findings carry the paper. First, a measured-manifold fine-tune recovers a model crippled by a fixed array signature from 32.0° back to 9.1° azimuth RMS (within 0.6° of a from-scratch oracle) using only 2000 captures. Second, sequential multi-site Kalman fusion reduces median horizontal error from 71.8 m (a fair per-snapshot azimuth-plus-elevation intersection) to 18.3 m. Third, bit depth is decisive only in a loud RF environment: under a 50 dB in-band blocker the 8-bit array produces gross bearing errors 88% of the time versus 2% for the 16-bit array — but an adversarial model with imperfect blocker cancellation, third-order intermodulation, and converter spurs shows that this margin is unlocked only by deep (≥40 dB) cancellation and nearly vanishes at the ~30 dB a practical canceller achieves. Fourth, for azimuth on a quiet single emitter the five-element ring outperforms the four-element square (~3° versus ~18–25°); the advantage is the fifth channel, not the ring shape, and reshaping four elements does not help. Every reported number is simulation on a modeled manifold; converting them to field numbers requires known-angle captures from the physical arrays, which the released toolchain ingests directly.

## 1. Introduction

Small unmanned aircraft are detected, classified, and localized through several modalities; the radio link between aircraft and operator is among the most useful, because it is present whenever the aircraft is under control and it carries a direction an interceptor can act on [16]. Passive RF direction finding on that link places a bearing on the emitter without radiating, and a network of bearings localizes it. The sensing quality then rests on two things: the receiver hardware, and the estimator that turns its samples into an angle.

Two receiver classes bracket the cost spectrum. The Epiq Sidekiq X4 provides four phase-coherent receive paths with 16-bit converters over 70 MHz–6 GHz [20]. The KrakenSDR provides five phase-coherent channels built from 8-bit RTL2832U tuners sharing one clock, at roughly two orders of magnitude lower cost, and has been used for real-world interference localization with a uniform circular array and MUSIC [14,15,21]. The natural question for a program that must field sensors at scale is what the extra cost of the 16-bit array actually buys, and whether the cheap array can be pushed to close the gap. This paper answers both through simulation, and—because a simulation that cannot see the difference between the two receivers would silently favor the cheaper one—it treats its own methodology adversarially.

**Contributions.** (i) A shared, impaired array-manifold simulation on which classical subspace methods [1,3,6], a direct neural classifier [12], and seven enhancements are measured against the Cramér–Rao bound [4], each with a planted-fault self-test. (ii) A two-platform model that adds ADC quantization, calibration residual, and phase noise, isolating what bit depth and geometry each contribute. (iii) An adversarial analysis of the dynamic-range result under a realistic analog front end. (iv) A link-budget conversion from accuracy to detection range. (v) A composition study of the stacked enhancements. (vi) A released, self-testing codebase with trained per-platform models and a real-capture ingestion path.

## 2. Background and related work

Subspace DF descends from MUSIC [1] and ESPRIT [2], with root-MUSIC giving a gridless polynomial form for uniform linear arrays [3] and the Cramér–Rao bound setting the estimator-independent variance floor [4]. Coherent sources, as from multipath, collapse the signal subspace; spatial smoothing [5] and its forward–backward variant [6] restore rank, but forward–backward smoothing assumes a centrosymmetric array, a point that matters in Section 5. Sparse geometries — minimum-redundancy [7], nested [8], and coprime [9] arrays — synthesize a difference coarray whose virtual aperture resolves more sources than physical elements. The matrix pencil method estimates angles from a single snapshot [10].

Data-driven DF has grown quickly. Deep classifiers tolerate array imperfections that degrade model-based methods [12]; hybrid model-data estimators such as SubspaceNet learn a surrogate covariance inside a differentiable root-MUSIC, keeping subspace interpretability while coping with coherence, wideband signals, low SNR, and array mismatch [11]; model-based deep architectures augment the covariance or subspace stages directly [13]. For low-cost coherent arrays specifically, the KrakenSDR has been characterized as hardware [15] and used for field interference localization with a circular array [14]. Multi-sensor bearing fusion and bearings-only tracking are standard [17]. The receiver-side limits we invoke — quantization signal-to-noise ratio and spur-free dynamic range [18], and third-order intermodulation from front-end nonlinearity [19] — are textbook, but to our knowledge have not been folded into a like-for-like DF comparison of an 8-bit and a 16-bit coherent array for this application.

## 3. Signal model, methodology, and validation discipline

All studies share one array-manifold generator. A narrowband source at azimuth and elevation (az, el) produces the per-element steering vector for a given geometry; the received snapshot adds per-element gain error (0.5 dB), phase error (5°), nearest-neighbour mutual coupling, 2 mm position jitter, a coherent multipath ray in a third of draws, and complex Gaussian noise at an SNR drawn across 0–30 dB with 64–512 snapshots. Estimators consume the sample covariance; classical reads use MUSIC with forward–backward smoothing where the geometry permits it [1,6], and the neural reads use a covariance-feature classifier [12]. The Cramér–Rao bound is computed in the deterministic form [4] as the reference floor.

**Validation discipline.** Each study carries a planted fault: an injected error that was observed to turn the test red, so that a passing run is evidence rather than decoration. A method is credited as an improvement only when it beats what already exists — classical MUSIC with forward–backward smoothing, or the existing neural classifier — on held-out data, not a weakened baseline. For every metric that gates a claim we ask the cheapest way it could pass without the underlying outcome, and verify that path is closed; several defects described below were caught this way. The full suite runs end-to-end and is reproduced clean on each change.

Figure. Array geometries evaluated

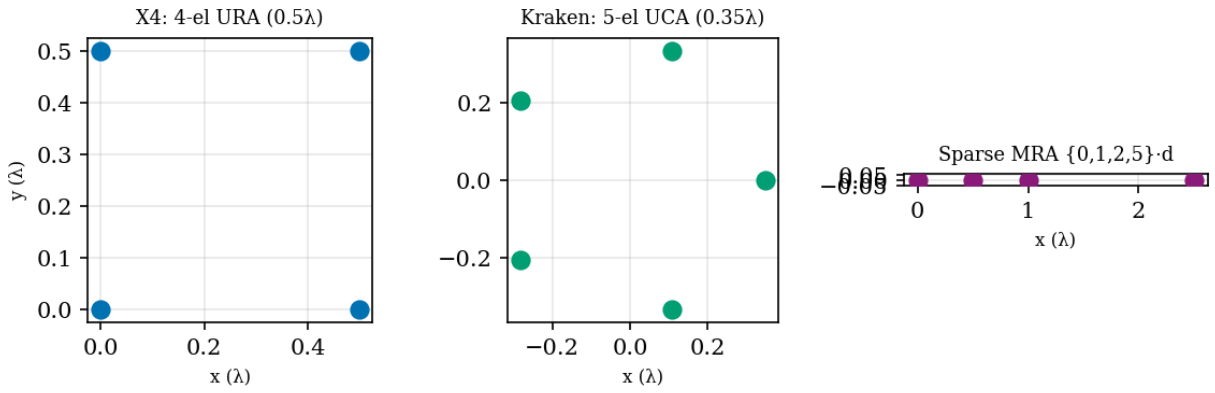



*Figure 1. Array geometries evaluated: the X4 four-element square (URA), the KrakenSDR five-element ring (UCA), and the four-element minimum-redundancy array whose difference coarray spans 11 virtual elements.*

## 4. Baseline direction finding and seven enhancements

The baseline established that a coherent four-channel array localizes a small drone by angle of arrival where the time-difference method fails on narrowband links. Seven enhancements were then measured on the shared manifold; the scorecard is Figure 2. Three pay off, one is an enabler, and three do not pay off on four channels — reported as findings, because the purpose was to find out.

**Figure 2. Honest scorecard after measurement -- three of seven do not pay off on 4 channels**

| | | | |
|---|---|---|---|
| 1 | Measured-manifold fine-tun | WIN | 7.9->32 deg degraded; recovered to 9.1 (oracle 8.5) |
| 2 | Sequential multi-site fusion | WIN | 72 m -> 18 m; conf. payoff NEES 12 -> 2 |
| 6 | Sparse MRA array | WIN | 4 elements resolve 4-5 sources via the difference coarray |
| 4 | Real-time inference | OK | sub-ms/fix; dwell-limited, not compute-limited |
| 3 | Learned-subspace denoiser | NO GAIN | barely beats raw MUSIC; classifier stays better |
| 5 | Gridless (root-MUSIC / MPM | NO GAIN | ~ grid at 4 elements; needs >4 elements |
| 7 | Wideband focusing | NO GAIN | no benefit on 4 elements (code correct) |

*Figure 2. Honest scorecard of the seven enhancements after measurement.*

### 4.1 Measured-manifold fine-tune (win).

A model trained on random impairments tolerates random spread but not a fixed per-array signature — the specific gain, phase, and coupling of one physical array. Modelling that signature degrades the deployed classifier from 7.9° on its nominal manifold to 32.0° on the real array. Fine-tuning on 2000 real captures recovers it to 9.1°, within 0.57° of a model trained from scratch on the real array (8.54°), while a control fine-tuned on the same quantity of nominal data stays at 26.7° — the recovery comes from the right data, not merely more training (Figure 3). An earlier version of this test mistakenly used higher-variance random impairments as the real array and showed no benefit; the premise is a fixed signature, and the test was rebuilt accordingly.

Figure 3. Fixed array signature cripples the model;
a few real captures recover it to near-oracle

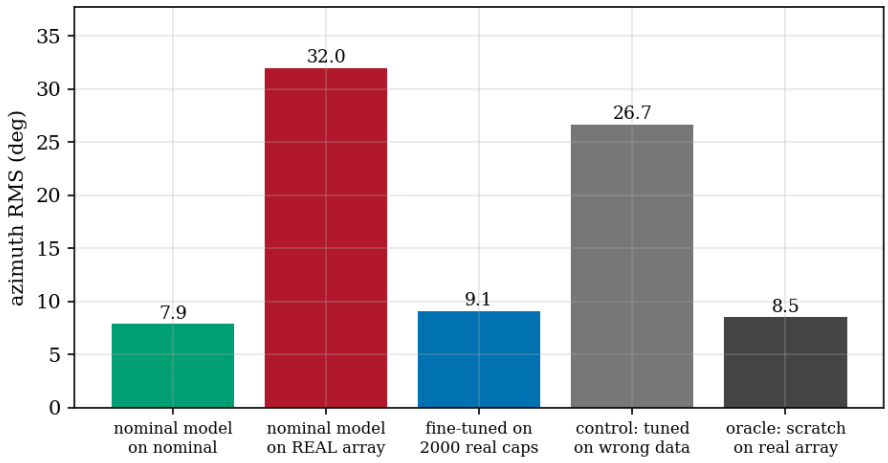



*Figure 3. A fixed array signature cripples the model; a few hundred to a few thousand real captures recover near-oracle accuracy.*

### 4.2 Sequential multi-site fusion (win).

Feeding the time series of calibrated bearings from several sites into an extended Kalman filter cuts median horizontal error from 71.8 m, a fair per-snapshot azimuth-plus-elevation intersection, to 18.3 m (Figure 4). The gain is temporal: the same per-snapshot intersection using elevation is already 71.8 m, so the filter's reduction is time, not the elevation dimension. Calibrated confidence is the sharper result — a filter that inflates the measurement variance of low-confidence bearings reaches

18.3 m and stays statistically consistent (normalized estimation error squared 2.0 against an ideal near 3), while one that trusts every bearing equally reaches 42.5 m and is inconsistent (11.8). A sign error in the measurement Jacobian diverges the filter, which the self-test catches.

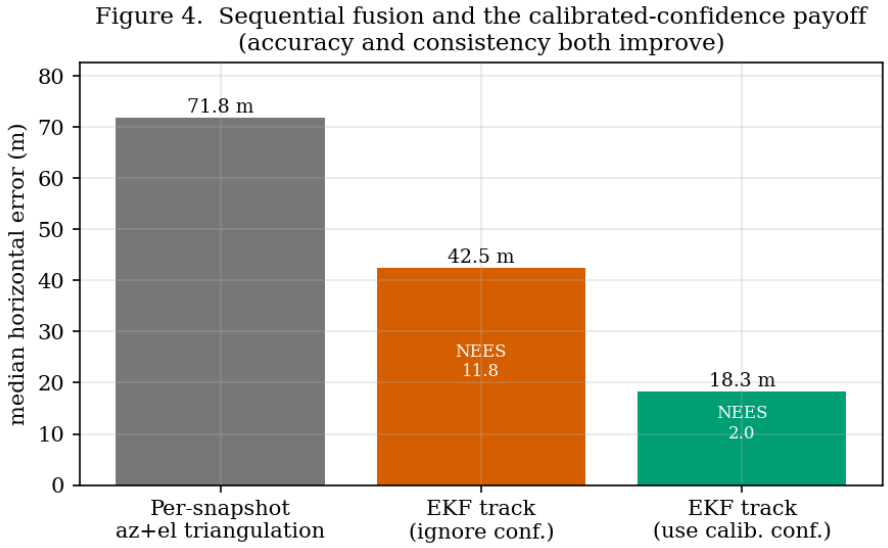



*Figure 4. Sequential fusion and the calibrated-confidence payoff.*

### 4.3 Sparse minimum-redundancy array (win).

Four elements at {0,1,2,5}·d synthesize an 11-element difference coarray (lags −5..+5) [7,8]. Coarray spatial-smoothing MUSIC on it resolves four sources 70% of the time and five 40% of the time, where a filled four-element line array never resolves four; the wider virtual aperture also lowers the single-source bound. The counter-metric: coarray processing assumes uncorrelated sources and collapses under coherence, so for a single drone it buys nothing — its value is separating several co-band emitters.

### 4.4 The enabler and the three that do not pay off.

Inference is sub-millisecond, so update rate is set by radio dwell, not compute. The three super-resolution and wideband ideas do not beat what already exists on four channels: a learned covariance denoiser of the SubspaceNet family [11] barely improves on classical MUSIC because the hard cases are geometric manifold ambiguity rather than covariance noise; gridless root-MUSIC ties the grid until the array grows past four elements; and wideband focusing shows no benefit at this aperture. A sign/subspace bug in the matrix pencil implementation [10] was found and fixed during this work (it degraded as elements were added); the corrected method still loses to MUSIC at four channels, so the conclusion survived the fix. Figure 5 shows the gridless result.

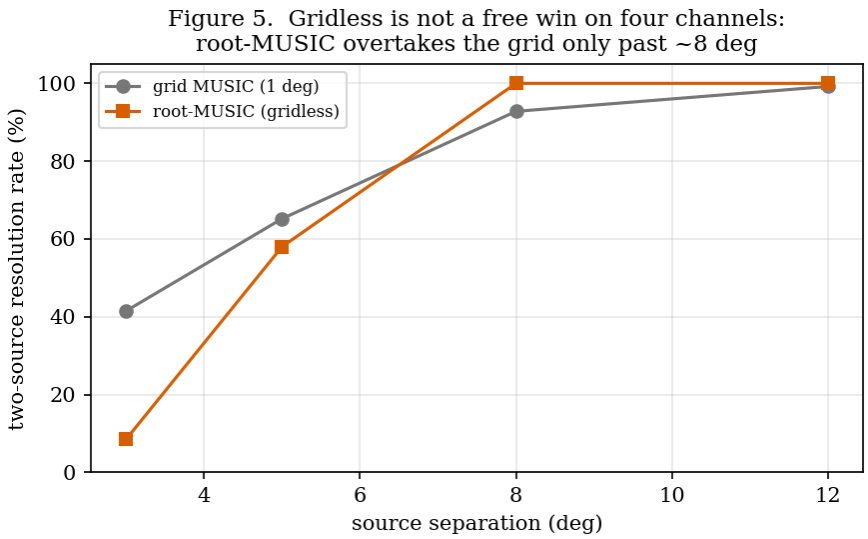



*Figure 5. Gridless estimation is not a free win on four channels.*

## 5. Two-platform model: what bit depth and geometry each contribute

The baseline manifold quantized nothing, so an 8-bit and a 16-bit receiver were indistinguishable on it. We add a uniform mid-rise quantizer with full scale set by the strongest in-band signal, a fixed per-element calibration residual, and LO phase noise, and compare the X4 (four-element square, 16-bit) against the Kraken (five-element ring, 8-bit) with everything else held equal.

**A defect caught here.** Forward–backward smoothing [6] is valid only for a centrosymmetric array. The square is centrosymmetric; the five-element ring is not. Applied to the ring, smoothing scrambled the signal subspace and produced 75%% gross errors at 20 dB SNR — which, run uncorrected, would have 'proven' the X4 dominates for a spurious reason. With smoothing disabled for the ring the figure falls to about 1%%. The estimator is now geometry-aware, guarded by a regression test.

**Dynamic range is the X4's real edge, under the idealized blocker.** With no strong signal the two bit depths tie; the converter is not the limit on a quiet drone. Under a 50 dB in-band blocker that sets the receiver gain and is then removed, the 8-bit array throws gross bearing errors 88% of the time while the 16-bit array holds at 2%, matching a transparent 24-bit control; the knee sits near the 45 dB the 8-bit quantization floor predicts [18] (Figure 6a). **For quiet-target azimuth the five-element ring wins** (~3° versus the square's 18–22°): more channels and a ring aperture simply beat four co-planar elements (Figure 6b). The largest software lever for the cheap array is calibration, not bit depth: shrinking the per-element phase residual from 25° to 5° moves single-source RMS from 7.3° to 3.1°; channel count and dwell help next.

Figure 6a.  Dynamic range is the X4's edge under the idealized blocker: a strong in-band signal blinds 8-bit; 16-bit rides over it

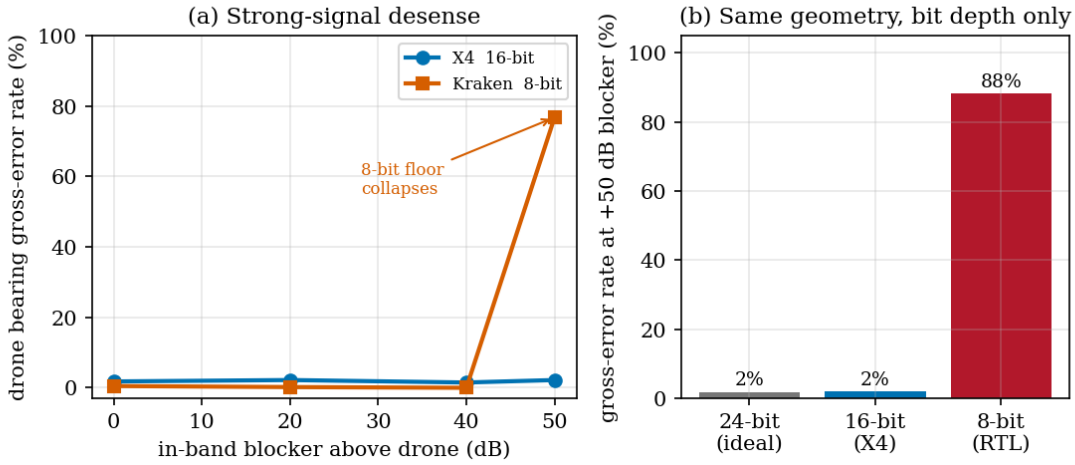



*Figure 6a. Dynamic range under the idealized (perfectly cancelled) blocker: 8-bit collapses, 16-bit rides over it and matches a transparent converter.*

**A correction we owe the record.** An earlier draft concluded from the ring-versus-square gap that a four-element array should be reshaped into a ring or line. Measured directly, that is wrong: among four-element layouts the compact square is the best azimuth geometry (0.4% gross), a four-element ring is worse (14%), and a line is far worse (72%, a cone and front/back ambiguity it cannot resolve). The X4's azimuth limit is its channel count, not its square shape; the lever is more channels, not a reshape.

## 6. Adversarial analysis, detection range, and composition

### 6.1 The dynamic-range win is conditional (adversarial).

The Section 5 blocker result assumed the blocker is perfectly cancelled, leaving only quantization noise. A real front end facing a strong blocker also suffers imperfect cancellation, third-order intermodulation that rises as the cube of the blocker power [19], and converter spurs at the spur-free dynamic range [18]. Adding these, the 16-bit advantage over 8-bit (gross-error-rate points) depends on cancellation depth: 1 points at 20 dB, 5 at 30 dB, 47 at 40 dB, and 88 at the near-perfect 120 dB the idealized model assumed (Figure 7). At the ~30 dB a practical adaptive or spatial canceller achieves on a four-element array, the 16-bit edge is a few points, not the tens the idealized model implied. The 16-bit converter is worth its cost for this task only in combination with good blocker cancellation; it is not automatic.

Figure. Adversarial realism: imp8's clean 16-bit win is conditional on near-perfect blocker cancellation

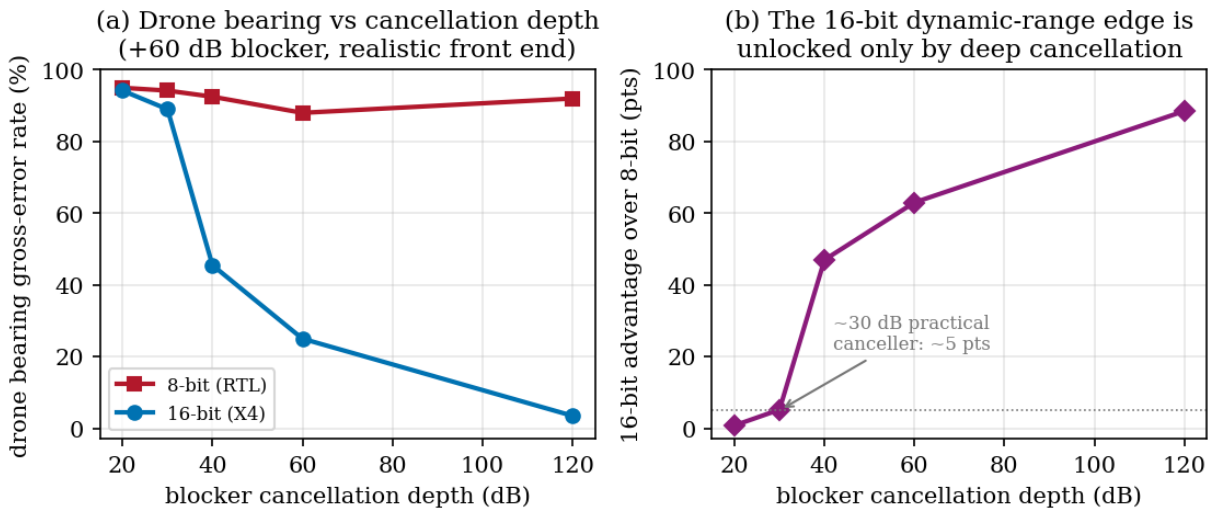



*Figure 7. Under a realistic analog front end, the 16-bit dynamic-range advantage is unlocked only by deep blocker cancellation and nearly vanishes at practical cancellation depth.*

### 6.2 From accuracy to detection range.

A link budget converts accuracy into the operator's question. With a consumer control/video-link EIRP of +20 dBm at 2.4 GHz, a log-distance path-loss exponent of 2.7 for a semi-cluttered low-altitude path, and per-platform noise figures of 4 dB (X4) and 7 dB (Kraken), the two platforms win on different axes (Figure 8). The X4's lower noise figure gives a longer coarse-bearing (track-loss) range — 3.5 km versus the Kraken's 2.2 km — while the four-element square's azimuth floor keeps its fine (≤10°) range near zero, and the five-element ring holds a fine bearing to 1.6 km. A 2 W (+33 dBm) link extends the ring's 10° range to 5.0 km. Free-space loss is the optimistic upper bound; real low-altitude paths are shorter.

Figure. Link-budget range: X4 holds a coarse bearing farther (lower NF); the Kraken ring holds a FINE bearing (geometry)

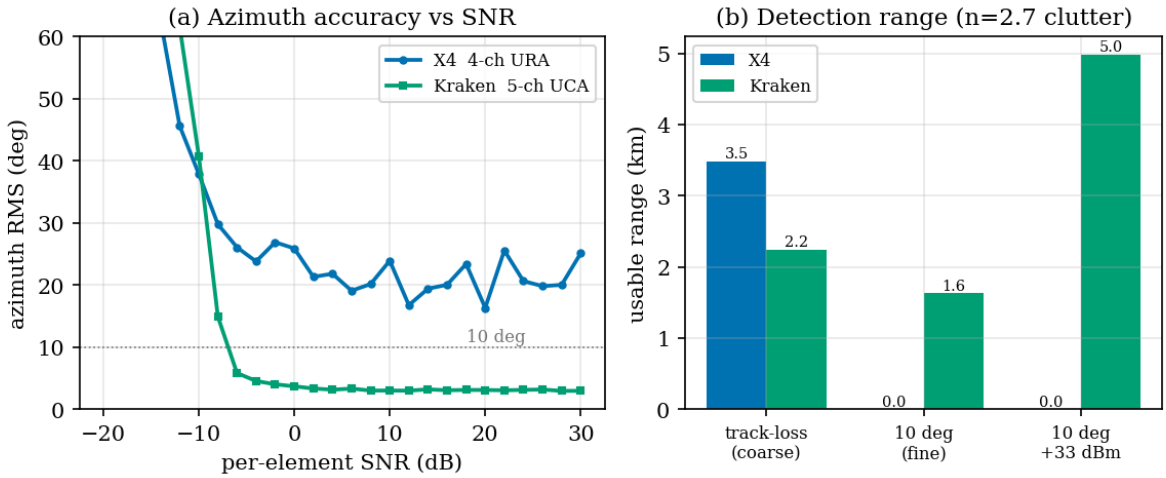



*Figure 8. Link-budget detection range. The X4 holds a coarse bearing farther; the ring holds a fine bearing.*

### 6.3 Do the wins compose?

Stacking the fine-tune and the fusion end-to-end on a single-drone localization, the gains compound but sub-multiplicatively (Figure 9). Median horizontal error falls from 177 m at baseline to 69 m with the fine-tune alone (2.6×), 21 m with fusion alone (8.6×), and 18 m stacked (10.0×). The stacked result beats either alone but is below the 22.1× naive product of the individual gains, because the filter's calibrated-confidence gross rejection already removes some of the ambiguity errors the fine-tune

reduces. The practical reading: do both — the fine-tune is cheap data and the fusion is pure software — and expect roughly a ten-fold, not a twenty-fold, improvement.

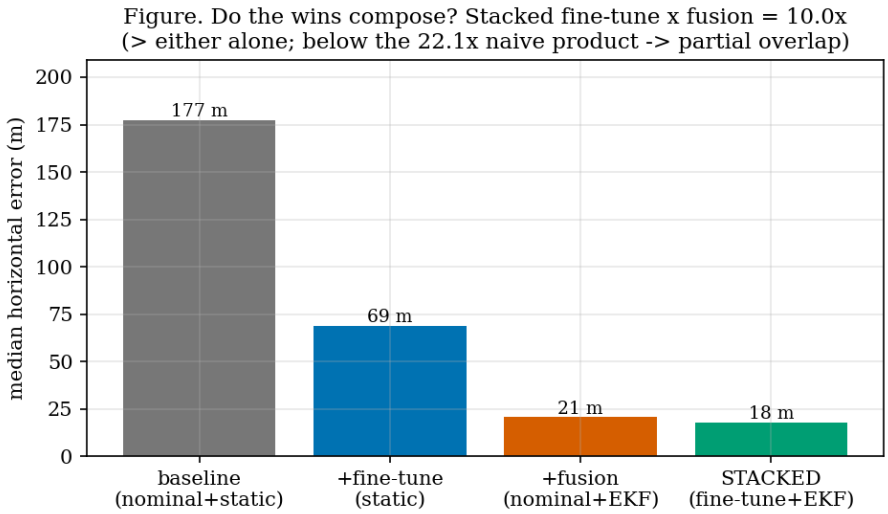



*Figure 9. The fine-tune and fusion gains compound, with partial overlap.*

### 6.4 Robustness: array-specificity and Cramér–Rao efficiency.

Two counter-metrics stress the claims above. First, how array-specific is a trained model? The per-platform networks were trained on a mild fixed signature. Tested on held-out arrays of increasing signature strength, the five-element ring model holds at its training accuracy for a mild held-out signature (4.9° in-distribution versus 5.5° at a 5° held-out spread) but degrades on a strong one (15.8° at 30°), so a per-array fine-tune is needed precisely when the real array's fixed signature is strong — which reconciles the fine-tune result of Section 4.1 with the mild-signature generalization seen elsewhere. The four-element square is ambiguity-limited (19.5°) and barely moves with the signature; its error is set by geometry, not calibration (Figure 10a). Second, classical MUSIC on the impaired array sits far above the deterministic Cramér–Rao bound computed on the ideal array — for the ring at 20 dB, 3.06° achieved against a 0.084° bound, a ratio near 36 (Figure 10b). No estimator beats the bound, and the large gap says the ring's accuracy is impairment- and calibration-limited rather than information-limited: the headroom is in calibration, consistent with Section 5.

Figure. Robustness: models are array-specific only for strong signatures; accuracy is impairment-limited, not information-limited

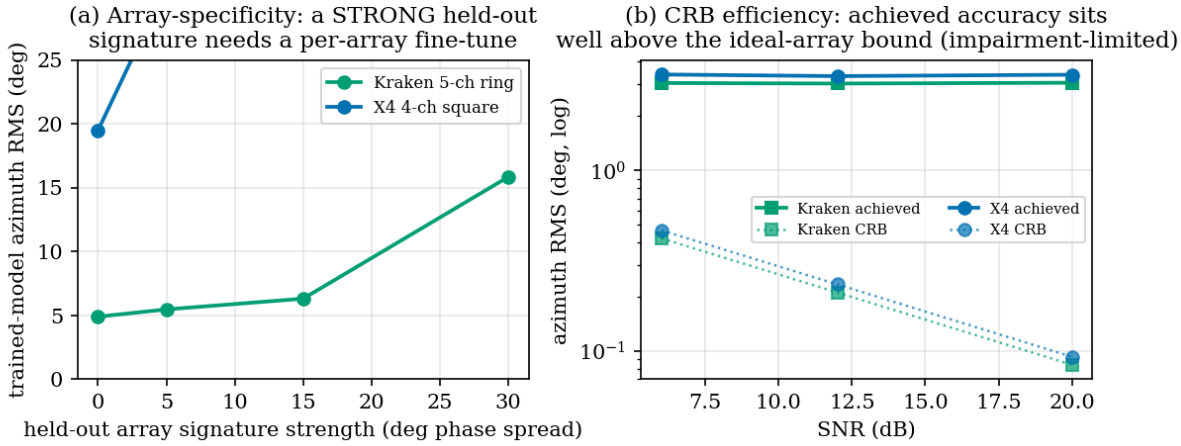



*Figure 10. (a) Trained-model error versus held-out array signature strength; a per-array fine-tune is needed only for strong signatures. (b) Achieved accuracy against the Cramér–Rao bound: impairment-limited, not information-limited.*

### 6.5 Realistic emitter and fusion sync robustness.

The fusion result of Section 4.2 assumed the sites share a clock. Injecting a per-site clock offset and sweeping it, the filtered horizontal error drifts only from 18.2 m at perfect sync to 22.6 m at a 2 s offset, against 69.5 m for per-snapshot triangulation — sub-second synchronization, trivial with GPS-disciplined timing, costs nothing, and the gain erodes only at multi-second offsets (Figure 11a). Site synchronization is not a practical limit on this fusion. A frequency-hopping or duty-cycled control link, which sits in the DF channel only part of the dwell, degrades DF gracefully rather than breaking it: the ring's core azimuth RMS holds near 3.1° from full duty down to 10% in-band duty, the loss appearing as a modest rise in the gross rate from fewer effective snapshots (Figure 11b).

Figure. The fusion win survives realistic site sync error; DF survives a duty-cycled hopping emitter

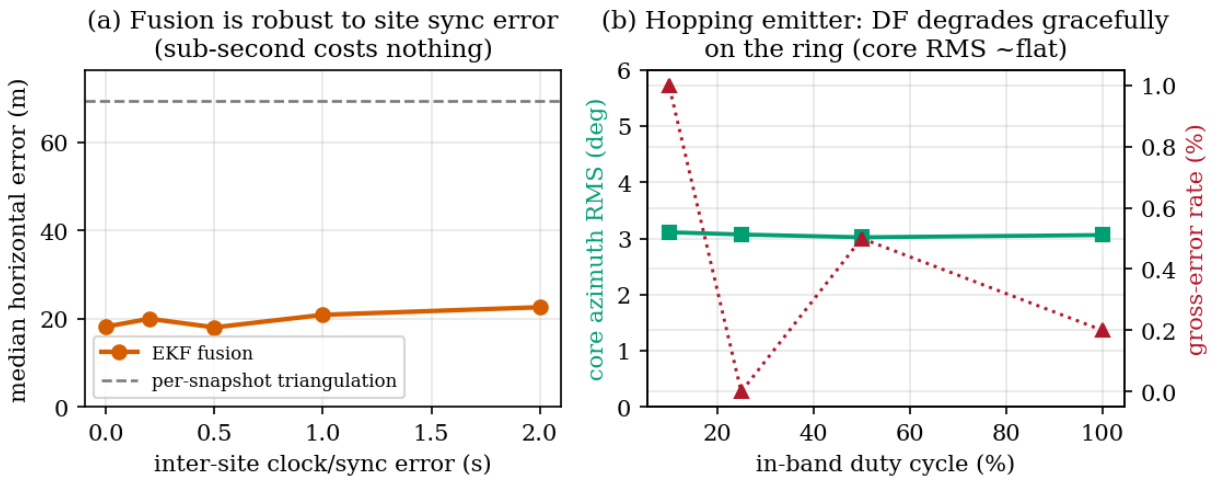



*Figure 11. (a) The fusion win survives realistic inter-site sync error. (b) DF survives a duty-cycled hopping emitter on the ring.*

### 6.6 A better decoder, and a hybrid that does not pay.

Almost all of the fine-tuned classifier's azimuth RMS comes from a small fraction of gross flips, not from its core accuracy. Reading the bearing as a circular soft-argmax over the whole softmax, rather than the single top bin, costs nothing and cuts RMS from 9.9 to 8.7 deg and the gross rate from 21% to 16% (the core holds near 4.6 deg). A per-fix hybrid that lets classical MUSIC referee the network's top-two azimuth candidates does NOT help on four channels, and we record it as a negative: MUSIC on the tripod flips more often than the network, so even when it only chooses between the net's own two peaks the blind referee is worse (9.2 deg) and a calibration-aware referee that knows the array's fixed signature merely ties it (8.5 deg). The azimuth ambiguity is geometric and shared across estimators; it is dissolved by more channels or by the temporal fusion below, not by a cleverer single-snapshot referee [11]. The soft-argmax read-out is wired into the real-capture path (captures.py), which loads the trained per-geometry model and decodes the field bearing with it, so a deployed capture carries this improvement rather than a plain argmax or classical MUSIC alone.

### 6.7 A mixed network of cheap and premium receivers.

Place several sites around the target, each a Kraken or an X4, and fuse their bearings [17]. In a quiet band the sharp five-element ring dominates: six Krakens reach 13 m against 35 m for six X4s, whose wide azimuth is dead weight. Under a strong blocker that blinds the 8-bit front ends, the answer splits on how fast a fix is needed. For a single-snapshot fix the cheap net collapses as more sites are jammed while two blocker-proof X4 anchors hold, and the mixed network overtakes pure-Kraken once the blocked fraction reaches 90%. But a 40-step EKF pulls the occasional clean Kraken bearing out of the jammed stream and keeps the cheap net ahead at every blocking level (crossover: never (the cheap net wins at every level)). The 16-bit dynamic range earns its hundred-fold cost in this network only for fast fixes in heavily contested spectrum; given seconds to fuse, density of cheap nodes wins (Figure 12).

Figure. Mixed RTL/Kraken + X4 network: the X4's dynamic range pays only for FAST fixes under heavy blocking (fast crossover 0.9); time fusion rescues the cheap net (EKF crossover None)

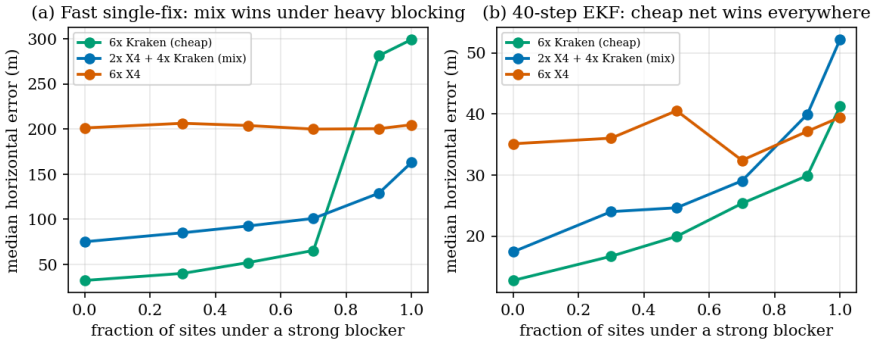



*Figure 12. Mixed RTL/Kraken + X4 network vs blocking fraction. (a) fast single fix: the X4 anchors win under heavy blocking. (b) 40-step EKF: temporal fusion keeps the cheap net ahead throughout.*

## 6.8 A network of phones.

A phone has one antenna and yields no bearing; a single phone is almost useless for location. A network of phones at known GPS positions still localizes the emitter two ways [23]. By received signal strength the swarm trilaterates the path loss with no time synchronization: median error falls from 91 m with five phones to 22 m with forty, set by shadowing (12 m at 2 dB, 37 m at 10 dB) rather than by detection, since a +20 dBm link is hearable to tens of kilometres. If the phones are GPS-time-synced to tens of nanoseconds, time-difference-of-arrival gives a sharper hyperbolic fix (11 m at 50 ns, degrading to 121 m at 500 ns), at the price of a coincident wideband capture of a known preamble that phones can rarely perform on an arbitrary link. A phone swarm is an area locator, not a bearing sensor, and only the network sees anything at all (Figure 13).

Figure. A phone swarm (single antenna, no bearing) localizes only as a NETWORK: RSS to tens of metres, GPS-synced TDOA sharper

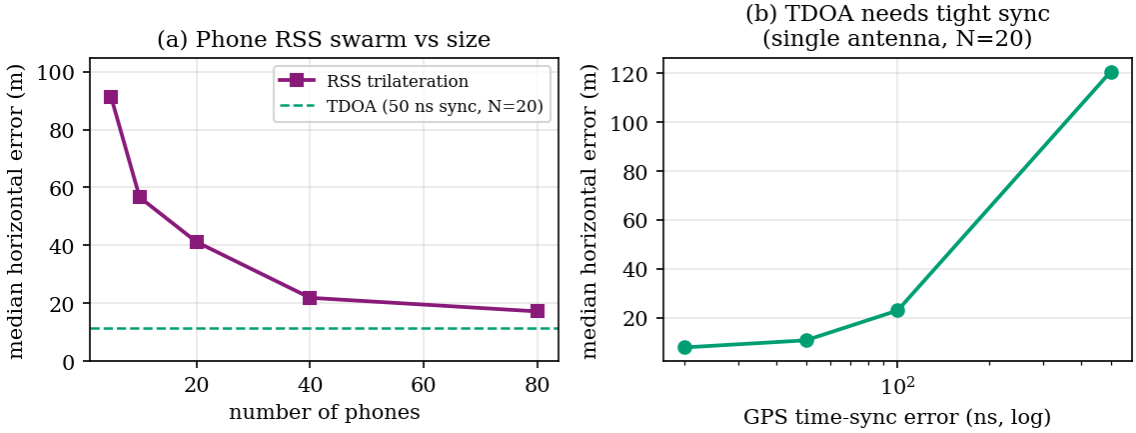



*Figure 13. A single-antenna phone swarm localizes only as a network: RSS trilateration to tens of metres (shadowing-limited), GPS-synced TDOA sharper.*

## 6.9 One estimator for every sensor.

Bearings and phone ranges fuse in one weighted estimator, each at its inverse-variance weight. A single bearing node cannot localize at all — one bearing is a line, not a point — yet adding forty phones yields a 33 m fix that beats the same forty phones alone (50 m), the lone bearing supplying the cross-range the swarm lacks. The converse caution is just as real: in a network already rich in sharp bearings (a three-ring Kraken net at 12 m), adding forty coarse RSS phones at full weight DRAGS the fix to 19 m, because received-strength ranging is biased and noisy; down-weighting the weaker sensor restores it to 15 m. Fuse every sensor, but trust each in proportion to what it actually measures (Figure 14).

Figure. One estimator, bearings + single-antenna phone RSS: phones fill a bearing-starved network and must be down-weighted in a bearing-rich one

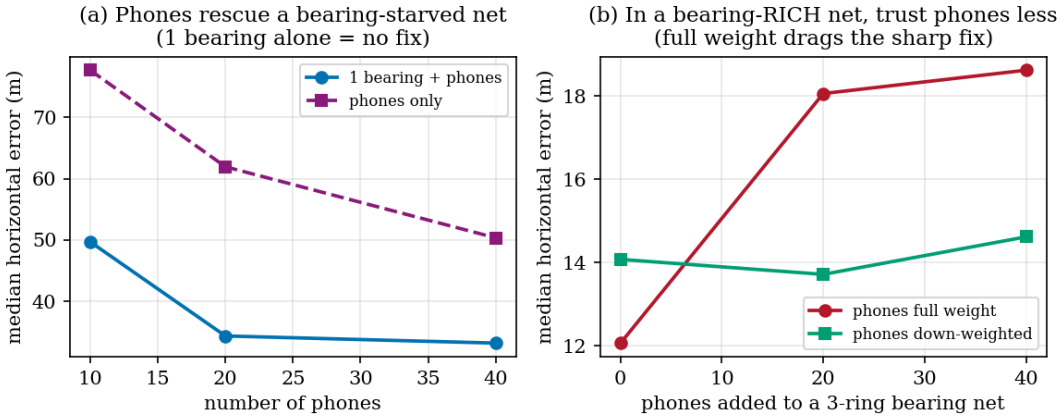



*Figure 14. Bearings + single-antenna phone RSS in one estimator: phones rescue a bearing-starved network and must be down-weighted in a bearing-rich one.*

## 6.10 Where to put the nodes, and what losing them costs.

Two operational questions close the network study: placement and failure. For six bearing nodes the geometry sets the fix. A full ring localizes to 14 m and covers 93% of the area under 30 m; a spread-out line reaches 18 m at 72% coverage; and a one-sided 60-degree cluster is the worst at 39 m and 42% — worse even than the line, because what helps is angular diversity around the target, not baseline length (the collinear-to-ring dilution is 1.3×). Nodes also fail. Dropping sites from the six-ring one at a time, the error degrades gracefully — 15 m at five, 17 at four, 22 at three — down to a two-bearing floor (43 m), then falls off a cliff at a single bearing (140 m, essentially no fix), because two bearings are the minimum that intersect to a point (Figure 15). The design reading: ring the target, and keep at least three nodes in view for margin above that floor.

Figure. Where to put the nodes and what node loss costs: a full ring beats a line or a one-sided cluster; the fix degrades gracefully to a 2-bearing floor

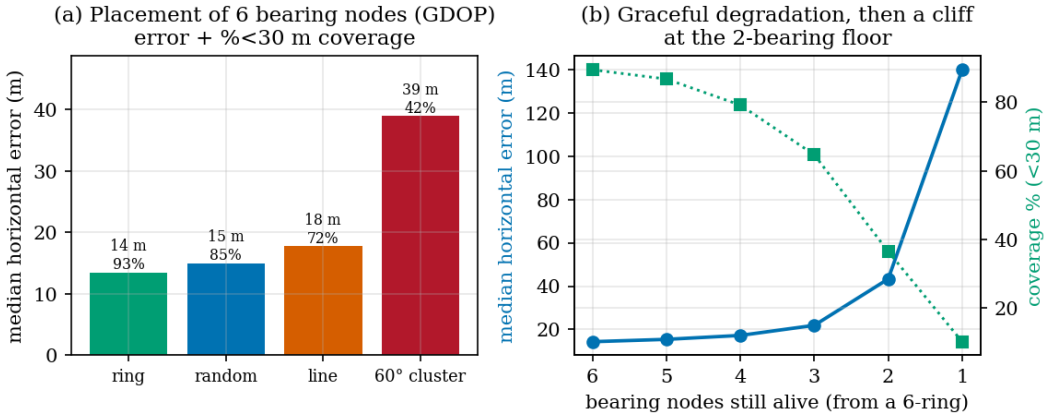



*Figure 15. (a) Placement of six bearing nodes: a ring beats a line and a one-sided cluster. (b) Graceful degradation under node loss to a two-bearing floor, then a cliff.*

## 6.11 Making the real-capture path survive real receivers.

Everything above runs on the synthetic manifold, and a model trained there meets defects on first contact the manifold never showed it: a DC term from local-oscillator leakage, per-element I/Q gain and phase imbalance, and hard clipping when the front end saturates. Injecting each into in-distribution captures and scoring through the fielded path, only one of the three moves the bearing. On the ring, I/Q imbalance to 30 degrees and clipping to a fifth of full scale leave azimuth RMS flat near 3 degrees; DC leakage at signal level, by contrast, drags raw MUSIC to 78 degrees as it locks onto the DC phantom. A DC-removal and blind I/Q-balancing step ahead of the estimators — now applied on every capture — returns it to 3.1 degrees and holds there across the sweep, for both the classical and the neural bearing, so no augmented retraining is needed before flight. A second

defect lived in the field path itself: forward-backward smoothing was applied to every geometry, and on the non-centrosymmetric tripod — the fine-tuned platform — that smoothing corrupts MUSIC rather than helping it (19 versus 4.8 degrees). The scorer now gates smoothing on the forward-backward validity condition — the index-reversal operator must line up with the array's geometric mirror, stricter than centrosymmetry alone (Figure 16a).

The array's own fixed signature is the last pre-field unknown, and the expensive one. A tripod carrying a strong per-element mismatch (about 3 dB and 40 degrees) reads 83 degrees RMS with 53% gross flips — effectively no bearing — if the manifold assumes an ideal array. Recording a handful of captures against a source at known azimuth lets the signature be estimated directly, as the dominant covariance eigenvector divided by the ideal steering, and folded into the manifold. That calibration brings the same array to 6.5 degrees at 0% gross — its geometry floor, a 13-fold gain from 8 beacon captures (Figure 16b). The estimate needs no mutual-coupling model, and a wrong signature leaves the array broken, so it is the correct one doing the work. The capture path consumes a beacon set or a surveyed vector directly. None of this is a field number; it is the software made ready so the first real captures are not wasted on a defect the simulation already knew how to remove.

Figure. Pre-field hardening: of the receiver defects only DC leakage moves the bearing and conditioning removes it; beacon calibration takes a miscalibrated array to its floor

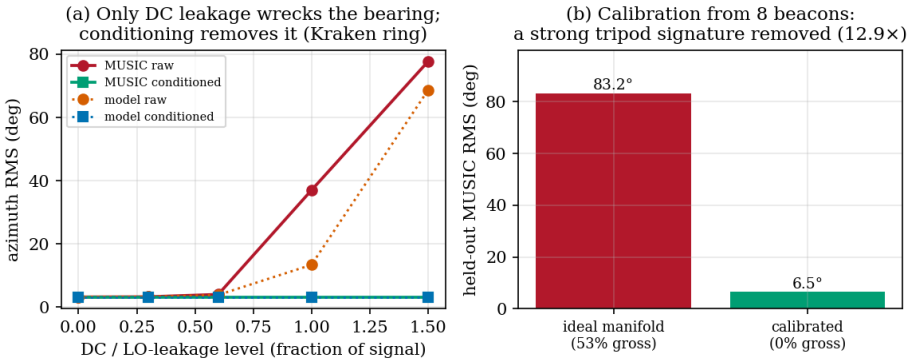



*Figure 16. (a) Of the receiver defects only DC leakage moves the bearing, and conditioning removes it. (b) Beacon calibration takes a strongly miscalibrated tripod to its geometry floor.*

### 6.12 Coherence recovery: direction finding on an array of independent receivers.

The arrays above assume phase-coherent channels. A low-cost alternative to a single multichannel receiver is an array of independent dongles — several RTL-SDRs, or HackRFs — but each carries its own oscillator phase, frequency offset, and sample-start timing, and those three unknowns sit on top of the geometry and destroy any phase-based bearing. A reference every channel hears at once (a switched noise source, a tone, or a chirp, split equally into all inputs) carries the same offsets, so they can be measured from it and subtracted: the common reference gives a rank-one covariance whose top eigenvector's per-element phase is exactly the correction, the same estimator as the array calibration of Section 6.11 pointed at a receiver reference. On a five-element ring at 2.4 GHz, independent per-channel oscillator phases leave MUSIC broken at 107 degrees RMS; one reference snapshot recovers it to 4.0 degrees, the coherent floor (4.0), while a deliberately wrong correction stays broken at 106 degrees — the measured reference phase is what fixes it, not any rotation.

The frequency offset sets the architecture. With the channels on a shared clock (a KrakenSDR, or HackRFs on a common CLKIN with the start trigger), the offset is zero and one reference snapshot holds for the session. Without a shared clock it does not: at 2.4 GHz a one-part-per-million inter-channel offset is 2.4 kHz, about half a cycle of phase walk across a 213-microsecond capture, and a one-shot correction that recovered 4.0 degrees at zero offset degrades to 46 degrees at one ppm. Re-estimating the reference per block, with the reference present throughout, tracks the drift back to 4.2 degrees. The practical reading: share a clock, or keep the reference live — either makes a pile of independent dongles into a usable DF array, and the recovery is validated here in simulation, not on hardware.

## 7. Discussion and deployment guidance

The measured picture is a division of labor, not a winner. The X4 is for contested spectrum: its 16-bit dynamic range and lower noise figure hold a bearing where a strong co-channel signal blinds an 8-bit front end — provided the array can cancel that blocker, which Section 6.1 shows is the real precondition. The KrakenSDR is a capable, inexpensive azimuth finder in a quiet

band, limited chiefly by calibration rather than by its converter, and its fifth channel gives it an azimuth edge over a four-element square. A layered deployment follows directly [16]: many cheap ring nodes for wide-area azimuth coverage, with 16-bit nodes where the RF environment is loud. On the algorithm side, the fine-tune and the fusion are the enhancements to field; the super-resolution and wideband methods should not be sold as improvements for four channels without more elements.

The network results sharpen the same lesson. A mixed receiver network is worth its cost only where it is needed: the premium 16-bit nodes buy fast fixes in a jammed band, while for anything the operator can watch for a few seconds a dense field of cheap rings plus temporal fusion wins (Section 6.7). Single-antenna phones carry no bearing, but a phone swarm is a real area locator by received strength (tens of metres) or GPS-synced timing, and a single bearing node fused with a phone swarm localizes where neither could alone (Sections 6.8–6.9). The uniting rule is to fuse every available sensor at its own measured accuracy, down-weighting the coarse ones rather than excluding or over-trusting them. A practical counter-UAS net is therefore heterogeneous by design: a few bearing arrays for precision, many cheap or opportunistic nodes for coverage and for the geometry that breaks ambiguity.

## 8. Limitations and threats to validity

Every number here is simulation on a modeled manifold, and the manifold, though impaired, is simpler than a real site. The mutual-coupling model is a simple exponential rather than an electromagnetic solution; the propagation model is log-distance rather than site-specific; the drone emitter is modeled as a complex-Gaussian source whose only waveform structure studied is a hopping duty cycle (Section 6.5), not the full modulation of a specific control link. The adversarial blocker model in Section 6.1, though it adds intermodulation and spurs, still treats the residual blocker as ideally subtracted apart from those terms, and its cancellation-depth axis is the honest way we expose that assumption rather than hide it. The calibration residual is modeled as a fixed complex per-element vector; thermal and frequency drift of that vector are not yet studied. The single hard dependency is empirical: converting any figure here to a field number requires known-angle four-channel captures from the physical X4 and a surveyed Kraken ring. The released toolchain ingests such captures through the same estimators, with a documented file format and a round-trip self-test, so the modeled tables become field tables without code change once the data exist.

## 9. Reproducibility

The code, the trained per-platform models (saved weights), the per-study result files, and a one-command reproducer live on the `lm-detector` branch under `tools/rf/doa_nn/improvements/`. Each study is a runnable module with a planted-fault self-test; `run_all.py` executes the suite and writes an aggregate summary. Trained weights were verified to reload and reproduce their reported azimuth RMS. No test, threshold, or baseline was weakened to obtain a result; where a method lost, the study reports the loss and the self-test validates only the method's internal correctness.

---

***Disclosure (result card).*** *Results produced by the author's CGDroneHunter simulation code (Python/NumPy/PyTorch, CPU), 29 October 2026. Status: simulation, not field-measured. All numbers reproduce from the cited JSON result files via* `run_all.py` *at the stated commit. Bibliographic details were verified against literature search where possible; readers should confirm exact page and issue numbers against each DOI before citing. No human steering of results; no measured field data is claimed.*

## References


[1] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," IEEE Trans. Antennas Propag., vol. 34, no. 3, pp. 276–280, 1986.

[2] R. Roy and T. Kailath, "ESPRIT—estimation of signal parameters via rotational invariance techniques," IEEE Trans. Acoust., Speech, Signal Process., vol. 37, no. 7, pp. 984–995, 1989.

[3] A. J. Barabell, "Improving the resolution performance of eigenstructure-based direction-finding algorithms," in Proc. IEEE ICASSP, 1983, pp. 336–339. (root-MUSIC)

[4] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and Cramér–Rao bound," IEEE Trans. Acoust., Speech, Signal Process., vol. 37, no. 5, pp. 720–741, 1989.

[5] T.-J. Shan, M. Wax, and T. Kailath, "On spatial smoothing for direction-of-arrival estimation of coherent signals," IEEE Trans. Acoust., Speech, Signal Process., vol. 33, no. 4, pp. 806–811, 1985.

[6] S. U. Pillai and B. H. Kwon, "Forward/backward spatial smoothing techniques for coherent signal identification," IEEE Trans. Acoust., Speech, Signal Process., vol. 37, no. 1, pp. 8–15, 1989.

[7] A. T. Moffet, "Minimum-redundancy linear arrays," IEEE Trans. Antennas Propag., vol. 16, no. 2, pp. 172–175, 1968.

[8] P. Pal and P. P. Vaidyanathan, "Nested arrays: A novel approach to array processing with enhanced degrees of freedom," IEEE Trans. Signal Process., vol. 58, no. 8, pp. 4167–4181, 2010.

[9] P. P. Vaidyanathan and P. Pal, "Sparse sensing with co-prime samplers and arrays," IEEE Trans. Signal Process., vol. 59, no. 2, pp. 573–586, 2011.

[10] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," IEEE Trans. Acoust., Speech, Signal Process., vol. 38, no. 5, pp. 814–824, 1990.

[11] D. H. Shmuel et al., "SubspaceNet: Deep learning-aided subspace methods for DoA estimation," IEEE Trans. Veh. Technol., 2023.

[12] Z.-M. Liu et al., "Direction-of-arrival estimation based on deep neural networks with robustness to array imperfections," IEEE Trans. Antennas Propag., vol. 66, no. 12, pp. 7315–7327, 2018.

[13] X.-X. Xu et al., "MD-DOA: A model-based deep learning DOA estimation architecture," IEEE Sensors J., 2024.

[14] L. Lehmann et al., "Real-world jammer localization using a low-cost array-based software-defined radio," in Proc. ION GNSS+ 2024, pp. 3715–3729.

[15] F. Taylor, E. Gattis, L. Trapani, D. Akos, S. Lo, T. Walter, and Y.-H. Chen, "Software defined radio for GNSS radio frequency interference localization," Sensors, vol. 24, no. 1, art. 72, 2024 (publ. 22 Dec. 2023), doi:10.3390/s24010072.

[16] I. Bisio et al., "Drone surveillance system—RF/WiFi-based drone detection, localization, and tracking: a survey," in Aviation Cybersecurity: Foundations, Principles, and Applications (ch. 3), IET, 2021.

[17] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, Estimation with Applications to Tracking and Navigation. Wiley, 2001.

[18] W. Kester, "Understanding SINAD, ENOB, SNR, THD, THD+N, and SFDR," Analog Devices Tutorial MT-003, 2009.

[19] B. Razavi, RF Microelectronics, 2nd ed. Prentice Hall, 2011. (IIP3, third-order intermodulation)

[20] Epiq Solutions, Sidekiq X4 Datasheet, rev. 1.012.16.24: four phase-coherent 16-bit receivers, 70 MHz–6 GHz.

[21] Realtek/Osmocom, RTL2832U: 8-bit I/Q on the USB bus; see Osmocom SDR documentation.

[22] H. Krim and M. Viberg, "Two decades of array signal processing research," IEEE Signal Process. Mag., vol. 13, no. 4, pp. 67–94, 1996.

[23] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero III, R. L. Moses, and N. S. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," IEEE Signal Process. Mag., vol. 22, no. 4, pp. 54–69, 2005.