Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Clutter-Aware Weak-Target Detection in Passive-Radar Range– Doppler Data Using Calibrated Recursive Complex-Valued Convolutive ICA

Hamzeh Ghahramani

hamze.ghahramani@gmail.com

Islamic Azad University, Bandar Lengeh

Research Article

Keywords: passive radar, weak-target detection, complex range–Doppler data, recursive complex-valued ICA, CFAR, clutter calibration, reproducibility

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11431862/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Clutter-Aware Weak-Target Detection in Passive-Radar Range–Doppler Data Using Calibrated Recursive Complex-Valued Convolutive ICA

Hamze Ghahramani

Department of Electrical Engineering, Islamic Azad University, Bandar-Lengeh Branch, Khamir Campus, Bandar Khamir, Iran

hamze.ghahramani@gmail.com | hamze.ghahramani@iau.ac.ir

Abstract

Weak-target detection in complex range–Doppler (RD) observations is challenging in heterogeneous clutter, particularly when raw reference and surveillance channels are unavailable. This study examines a clutter-aware detector combining frame-sequential recursive complex-valued convolutive independent component analysis (ICA) with interquartile-range (IQR) score calibration and maximum-rank fusion with a local CFAR-like baseline. Score orientation uses labeled calibration observations, whereas target-absent (H0) observations provide robust normalization and threshold estimates. Real-data evaluation uses only the processed complex RD measurements of the publicly available LSS-PR-1.0 dataset. In a pooled evaluation, the proposed H0-CG-OR-CVCICA detector achieved an area under the receiver operating characteristic curve (AUC) of 0.719, compared with 0.578 for the evaluated CA-CFAR-like baseline. Across 30 repeated file/group-level calibration/test splits, mean AUC was 0.714 ± 0.029 versus 0.570 ± 0.015, and mean detection probability at nominal false-alarm probability 0.01 was 0.260 ± 0.031 versus 0.132 ± 0.014. The proposed method had an empirical false-alarm probability of 0.015 and a reported processing time of 1.320 ms per frame in the tested Python implementation. Separate synthetic in-phase/quadrature (IQ) experiments were used exclusively to examine extensive cancellation algorithm (ECA)-type baselines, not for direct comparison on LSS-PR-1.0. Performance varied across target categories, and the two independent H0 clutter recordings limit conclusions about generalization to new recording conditions. The results support the usefulness of calibrated component-derived scores under the evaluated offline RD-domain protocol, without establishing end-to-end streaming operation or superiority over untested CFAR variants.

Keywords: passive radar; weak-target detection; complex range–Doppler data; recursive complex-valued ICA; CFAR; clutter calibration; reproducibility

1 Introduction

Passive radar exploits non-cooperative illuminators of opportunity and receives target echoes without transmitting a dedicated radar waveform. This operating principle makes passive radar attractive for covert sensing, spectrum-efficient surveillance, and persistent monitoring. However, passive reception also creates a difficult detection problem because the surveillance observation is a mixture of direct-path leakage, stationary and nonstationary clutter, multipath components, residual interference, thermal noise, and weak target echoes.

The problem becomes especially difficult for low signal-to-clutter ratio (SCR) targets. Small vessels, UAVs, helicopters, and other low-slow-small (LSS) targets may occupy only a compact region in the RD plane and can be partially masked by clutter or multipath-like residuals. Conventional passive-radar cancellation methods, including ECA, sliding ECA, least-squares suppression, multistage cancellation, STAP, sparse recovery, and group-sparse regression, provide strong baselines but rely on assumptions about raw reference channels, spatial calibration, dictionary construction, or target sparsity [1]-[8].

Complex-valued blind source separation (BSS) and independent component analysis (ICA) provide a complementary view. Instead of explicitly reconstructing a disturbance dictionary, BSS attempts to separate latent statistically independent components from complex observations. This is physically compatible with radar because phase and quadrature components are essential for Doppler processing. Convolutive ICA further represents each recovered component as a filtered combination of observations, making it suitable for mixture structures with memory, leakage, or local Doppler spread [9], [10].

An important distinction is that a recursively updated ICA separator does not automatically provide a calibrated detector or a fully causal end-to-end streaming system. Target-present calibration labels can be necessary to orient the score, while target-absent examples determine the robust normalization and false-alarm thresholds. The present study evaluates this processing chain offline using complex RD observations.

This study investigates H0-CG-OR-CVCICA, a clutter-aware detector incorporating frame-sequential recursive complex-valued convolutive ICA, supervised score orientation, H0-based median/IQR normalization, and CA-CFAR-assisted maximum-rank fusion. The term online recursive describes its separation-stage updates, not a demonstrated end-to-end real-time deployment. A separate controlled IQ simulation is used exclusively for ECA-type cancellation comparisons.

The contributions are: (1) a recursive complex-valued convolutive ICA separation stage for local passive-radar RD inputs; (2) an RD mixture construction suited to the available LSS-PR-1.0 observations; (3) calibration combining label-informed score orientation with H0-based median/IQR normalization and threshold estimation; (4) CA-CFAR-assisted maximum-rank fusion; and (5) offline pooled and repeated file/group-level evaluation with explicit limitations on clutter-recording independence and end-to-end streaming.

2 Related work

Passive radar clutter suppression and weak-target detection have been extensively studied through reference-channel cancellation, least-squares disturbance removal, STAP, sparse recovery, and threshold-setting strategies. ECA and sliding ECA remove direct-path and multipath replicas by projecting surveillance data onto delay-Doppler dictionaries generated from the reference signal [1], [2]. Least-squares and multistage methods provide related structured cancellation strategies [3], [4]. STAP exploits space-time covariance structure, but it generally requires calibrated multichannel array data [5]. Sparse and group-sparse methods formulate the detection problem through structured recovery and have shown strong performance for passive-radar masking and interference-cancellation problems [6]-[8].

Prior radar-oriented convolutive complex-valued ICA work motivates the choice of complex-valued separation [9], while conventional methods cover disturbance cancellation, spatial filtering and sparse recovery [1–8]. The present experiments examine the combination of recursive state updates and calibrated detection scores. The contribution is evaluated only within the specific experimental protocols described below; the study does not establish that the method outperforms untested CFAR variants.

The LSS-PR-1.0 benchmark provides processed complex RD records for passenger ships, speedboats, helicopters and unmanned aerial vehicles (UAVs), together with two independent sea-clutter recordings [11, 12]. Its public RD release supports the present detector evaluation, but the raw reference/surveillance IQ streams required for ECA are not available in the release used here.

3 RD-domain problem definition

The real-data experiment uses the LSS-PR-1.0 V3 complex RD release. The dataset contains target-present RD records for speedboats, passenger ships, helicopters, and UAVs, together with low- and high-sea-state clutter recordings. Only the complex RD observations are used in detector scoring; accompanying trajectory metadata are not used as detector input.

For each target-present record i, the complex RD cube is denoted by $\mathcal{Z}^{(i)} \in \mathbb{C}^{5\times 400\times K_i}$, where the first dimension contains five local range cells, the second dimension contains $D = 400$ Doppler bins, and K_i is the number of frames. Let $z_{r,d,k}^{(i)}$ denote its complex entry at local range index r, Doppler index d, and frame k. The central local range cell is the cell under test (CUT), and the outermost local range cells form the clutter-reference observation:

$$x_{\mathrm{cut},k}[d] = z_{3,d,k}^{(i)},$$
$$x_{\mathrm{ref},k}[d] = 1/2\left(z_{1,d,k}^{(i)} + z_{5,d,k}^{(i)}\right).$$

The corresponding two-channel complex observation is

$$\mathbf{X}_k = \begin{bmatrix} \mathbf{x}_{\mathrm{cut},k}^{\mathrm{T}} \\ \mathbf{x}_{\mathrm{ref},k}^{\mathrm{T}} \end{bmatrix} \in \mathbb{C}^{2\times D}, \qquad D = 400.$$

For target-absent trials, five-cell patches are extracted from the sea-clutter RD cubes $\mathcal{Z}_0^{(j)} \in \mathbb{C}^{101\times D\times K_j}$, where j indexes the source sea-clutter record and K_j is its number of frames. The resulting clutter-only observations are $\mathbf{X}_{0,q,k} \in \mathbb{C}^{2\times D}$, where q indexes the local range patch. The two detection hypotheses are

For H1, 43 source files provide 2383 observations: speedboat (12 files, 548 frames), passenger ship (11 files, 788 frames), helicopter (11 files, 668 frames), and UAV (9 files, 379 frames). These counts refer to available frame-level observations, not statistically independent acquisitions.

For H0, two sea-clutter files provide 1040 observations. The high-sea-state file contains 43 frames and contributes 860 five-cell patches; the low-sea-state file contains 9 frames and contributes 180 patches. For each frame, 20 non-overlapping five-cell range patches begin at zero-based range indices 0, 5, 10, …, 95. The middle cell of each patch is the CUT, and the two outer cells form its reference. Although H1 files are disjoint between split partitions, H0 patches from the same source recording may occur in both partitions. Matching target/clutter acquisition conditions has not been established.

$$H_0: \quad \mathbf{X}_k = \mathbf{C}_k + \mathbf{N}_k,$$
$$H_1: \quad \mathbf{X}_k = \mathbf{C}_k + \mathbf{T}_k + \mathbf{N}_k.$$

Here, $\mathbf{C}_k$, $\mathbf{T}_k$, and $\mathbf{N}_k$ denote the clutter and residual interference, target-related contribution, and noise, respectively, and have the same dimensions as $\mathbf{X}_k$. The binary decision is $\delta_k = 1$ when the calibrated detection statistic exceeds its H_0-derived threshold, and $\delta_k = 0$ otherwise.

4 IQ-domain simulation for ECA baselines

Because the LSS-PR-1.0 release used here does not provide raw reference and surveillance IQ channels, ECA and RLS-type baselines are evaluated only in a controlled simulation. Let $s[n]$ denote the illuminator complex baseband sequence, and let n denote the discrete-time sample index. The simulated reference channel is

$$r[n] = s[n] + v_r[n].$$

The simulated surveillance channel is

$$x[n] = \alpha_0 s[n] + \sum_{\ell=1}^{L_c} \alpha_\ell \, s[n-\tau_\ell]\exp(j2\pi f_\ell n T_s)$$
$$+\beta s[n-\tau_t]\exp(j2\pi f_t n T_s) + w[n].$$

Here, $v_r[n]$ and $w[n]$ denote reference-channel and surveillance-channel noise, respectively; α_0 is the direct-path coefficient; L_c is the number of modeled clutter/multipath components; α_ℓ, τ_ℓ, and f_ℓ are their complex coefficients, sample delays, and Doppler frequencies; and β, τ_t, and f_t are the corresponding target-related coefficient, delay, and Doppler frequency. T_s denotes the sampling interval and $j = \sqrt{-1}$. The simulation models $s[n]$ as a complex white Gaussian baseband illuminator for controlled algorithm validation. This simplification is explicitly stated and is not claimed to represent a specific broadcast waveform. The IQ-domain track evaluates cancellation depth and runtime for classical ECA, reduced-dictionary ECA, and block-wise reduced ECA.

5 Proposed detector

The proposed detector has five major stages: local RD mixture construction, lag embedding, online recursive complex-valued convolutive separation, labeled-calibration score orientation and H0-based median/IQR normalization, and CA-CFAR-assisted maximum-rank fusion.

For an RD frame $\mathbf{X}_k \in \mathbb{C}^{M\times D}$, a lag-embedding operator $\mathcal{L}_{L_w}(\cdot)$, which collects L_w adjacent Doppler-bin observations, forms $\mathbf{U}_k = \mathcal{L}_{L_w}(\mathbf{X}_k) \in \mathbb{C}^{P\times T}$, where M is the number of input channels, L_w is the embedding/filter length, $P = ML_w$, and $T = D - L_w + 1$ is the number of embedded Doppler samples per frame. For the LSS-PR-1.0 RD experiment, $M = 2$ and $D = 400$. Lag embedding models local Doppler-bin dependence without implying access to raw propagation-time IQ samples.

Recursive covariance estimation is performed as

$$\mathbf{R}_k = \lambda_R \mathbf{R}_{k-1} + \frac{1-\lambda_R}{T}\mathbf{U}_k\mathbf{U}_k^{\mathrm{H}}.$$

The covariance estimate $\mathbf{R}_k \in \mathbb{C}^{P\times P}$ uses the previous frame's covariance $\mathbf{R}_{k-1}$; $0 \le \lambda_R < 1$ is its exponential forgetting factor, and $(\cdot)^{\mathrm{H}}$ denotes Hermitian transpose. Diagonal regularization gives

$$\widetilde{\mathbf{R}}_k = \mathbf{R}_k + \varepsilon \mathbf{I}_P.$$

Here, $\varepsilon > 0$ is a small regularization constant and $\mathbf{I}_P$ is the $P \times P$ identity matrix. With the eigendecomposition $\widetilde{\mathbf{R}}_k = \mathbf{E}_k\mathbf{D}_k\mathbf{E}_k^{\mathrm{H}}$, where $\mathbf{E}_k$ contains eigenvectors and diagonal $\mathbf{D}_k$ contains eigenvalues, the whitening matrix is $\mathbf{V}_k = \mathbf{D}_k^{-1/2}\mathbf{E}_k^{\mathrm{H}}$ and the whitened observations are $\mathbf{Z}_k = \mathbf{V}_k\mathbf{U}_k$. Whitening may be updated once every K_V frames, where K_V is the whitening-update period.

The complex demixing matrix $\mathbf{W}_k \in \mathbb{C}^{P\times P}$ yields separated components $\mathbf{Y}_k = \mathbf{W}_k\mathbf{Z}_k$. The elementwise score/nonlinearity is $\varphi(\mathbf{Y}_k) = \mathbf{Y}_k/(|\mathbf{Y}_k| + \rho)$, where $\rho > 0$ is a denominator stabilizer, absolute value and division are elementwise, and $\varphi(\cdot)$ is applied elementwise. The natural-gradient update uses the step size $\mu_k > 0$:

$$\Delta\mathbf{W}_k = \left(\mathbf{I}_P - \frac{\varphi(\mathbf{Y}_k)\mathbf{Y}_k^{\mathrm{H}}}{T}\right)\mathbf{W}_k,$$
$$\mathbf{W}_{k+1}^{\mathrm{pre}} = \mathbf{W}_k + \mu_k\Delta\mathbf{W}_k.$$

To reduce component collapse, symmetric decorrelation is then applied to the pre-projection update as $\mathbf{W}_{k+1} = (\mathbf{W}_{k+1}^{\mathrm{pre}}(\mathbf{W}_{k+1}^{\mathrm{pre}})^{\mathrm{H}})^{-1/2}\mathbf{W}_{k+1}^{\mathrm{pre}}$, where $(\cdot)^{-1/2}$ denotes the inverse Hermitian square root.

The raw OR-CVCICA detection score $s_{\mathrm{or},k}$ summarizes component-level concentration, non-Gaussianity, reference decorrelation and temporal consistency. Its orientation is determined from labeled calibration observations: $a_{\mathrm{cal}} = -1$ if the raw calibration area under the receiver operating characteristic curve (AUC) is below 0.5, and $a_{\mathrm{cal}} = +1$ otherwise. Let $\mathcal{S}_{0,\mathrm{or}} = \{a_{\mathrm{cal}}s_{\mathrm{or},i} : i \in \mathcal{C}_0\}$ be the oriented scores from the H_0 calibration subset $\mathcal{C}_0$. The calibrated score is

$$s_{\mathrm{cal},k} = \frac{a_{\mathrm{cal}}s_{\mathrm{or},k} - m_0}{q_0},$$
$$m_0 = \mathrm{median}(\mathcal{S}_{0,\mathrm{or}}), \qquad q_0 = \max\{Q_{0.75}(\mathcal{S}_{0,\mathrm{or}}) - Q_{0.25}(\mathcal{S}_{0,\mathrm{or}}), \varepsilon_{\mathrm{score}}\}.$$

Here, $Q_p(\cdot)$ is the empirical p-quantile, q_0 is the interquartile range (IQR) of the oriented H_0 calibration scores, and $\varepsilon_{\mathrm{score}} > 0$ is the numerical lower bound used to prevent division by zero. Score orientation uses both H_0 and H_1 calibration labels; the robust location and IQR are estimated from oriented H_0 scores only. This is consistent with the v2 ScoreCalibrator implementation.

The calibrated component-derived score and the local CA-CFAR-like score $s_{\mathrm{cfar},k}$ are converted into normalized empirical ranks $r_{\mathrm{or},k}$ and $r_{\mathrm{cfar},k}$. The pooled score fusion is defined as the maximum of these two normalized ranks:

$$\mathcal{T}_k = \max\{r_{\mathrm{or},k}, r_{\mathrm{cfar},k}\},$$
$$r_{\mathrm{or},k} = \mathrm{rank}_{[0,1]}(s_{\mathrm{cal},k}), \qquad r_{\mathrm{cfar},k} = \mathrm{rank}_{[0,1]}(s_{\mathrm{cfar},k}).$$

The empirical $\mathrm{rank}_{[0,1]}$ transform in the pooled v2 code sorts the full score vector and assigns ranks from zero to one. No independent weight g_k or gate threshold is used in this maximum-rank rule. The detection decision is $\delta_k = \mathbf{1}\{\mathcal{T}_k > \eta(p_{\mathrm{fa}})\}$, where η is the empirical $(1-p_{\mathrm{fa}})$-quantile of H_0 fused scores. The pooled rank transformation is offline and uses all available scores; the repeated validation uses separate calibration and test partitions.

Recursive-state causality. Within the specified frame-processing order, covariance, whitening and natural-gradient demixing use the current complex RD observation and previously retained states; these update routines do not explicitly use future frames. The lag operator acts over Doppler bins of an already acquired frame, not the raw IQ time series. Separation-stage sequential processing must therefore be distinguished from the full downstream calibration, which is offline in the reported experiments.

6 Computational analysis

With M channels, lag length L_w, P = M L_w, and T = D - L_w + 1 samples per frame, recursive covariance estimation requires O(P^2 T). Full whitening by eigendecomposition requires O(P^3), but if it is updated every K_V frames, the amortized whitening cost is O(P^3/K_V). The demixing and natural-gradient update require O(P^2 T). Thus, the dominant per-frame complexity is

```
O(P^2 T) + O(P^3/K_V) = O(M^2 L_w^2(D-L_w+1)) + O(M^3 L_w^3/K_V).
```

The arithmetic cost is moderate for the tested short-lag configurations in the available Python implementation. Stability mechanisms include regularization, bounded learning rates, covariance forgetting and symmetric decorrelation. These measurements alone do not demonstrate end-to-end real-time behavior, since data access, calibration, rank fitting and decision threshold fitting are distinct processing operations.

The complexity sweep comprises L_w ∈ {2,3,5,7,9} and whitening-update intervals K_V ∈ {1,2,5,10}, evaluating 20 configurations. Its AUC and runtime results characterize a configuration trade-off within the tested offline data-processing setup rather than latency of an end-to-end operational radar deployment.

7 Experimental setup

The real-data track uses processed complex LSS-PR-1.0 RD observations. Baselines comprise raw CUT energy, CUT/reference ratio, one-dimensional CA-CFAR-like detection, SVD/whitening, batch cICA-lite, OR-CVCICA-v1, OR-CVCICA-v2 raw and calibrated, and the proposed score-fusion detector. The underlying positive and negative source recordings and sampling rules are described in Section 3 and Appendix A.

The separate synthetic IQ-domain track evaluates ECA-type cancellation baselines when simulated raw reference/surveillance channels are available. It reports cancellation ratio, runtime and dictionary size for classical batch, reduced-dictionary and block-wise reduced ECA. These simulations are not direct detection comparisons on LSS-PR-1.0.

Reported criteria are ROC AUC, detection probability (Pd) at nominal false-alarm probabilities (Pfa) 0.1 and 0.01, empirical Pfa, class-specific AUC and measured processing time. Complexity-sweep operating-point results are distinct from pooled and repeated split-fitted estimates; the experiments were conducted offline.

Baseline definitions. Before scoring, the two complex RD channels are separately centered and standardized to approximately unit average power. Raw CUT energy is the maximum over Doppler bins of the squared normalized CUT magnitude. CUT/reference ratio is the maximum normalized CUT power divided by the mean normalized reference-channel power, with a 10^{-12} denominator stabilizer. Independent normalization makes the latter denominator approximately unity, explaining the equal rounded results and reported thresholds of those two simple methods in Table II.

The CA-CFAR-like score is one-dimensional along 400 Doppler bins of the normalized CUT signal. For an interior candidate bin, its 57-bin neighborhood contains the candidate CUT, four guard bins on either side and 24 reference/training bins on each side (48 training bins total). The mean training-cell power excludes the CUT and guard cells. At Doppler boundaries the window is truncated without circular wrapping, and bins with fewer than four training cells are ignored. The frame score is the maximum valid ratio between CUT power and mean local training power, with a 10^{-12} denominator stabilizer. This is not a two-dimensional RD CFAR detector; the separate range-reference channel is not used to compute this baseline.

False-alarm thresholds are selected by the empirical (1−Pfa) quantile of H0 scores in the relevant calibration population. The pooled results use dataset-level scores and thresholds; the repeated validation reports calibration-partition fitting before application to test scores. Unlike conventional analytical CA-CFAR, no parametric exponential-clutter scale factor is used to set the reported frame-level threshold. Fully specified 2D CA-CFAR and ordered-statistics CFAR (OS-CFAR) were not tested.

Processing order and offline scope. The pooled experiment processes source files in filename order and target frames in increasing frame index; multiple clutter range patches are processed within each clutter frame. Recursive separation states may persist across recordings. Pooled score orientation and rank normalization use the pooled dataset and are therefore noncausal at the dataset level. The repeated file/group-level analysis uses calibration/test partitions and is also an offline evaluation, rather than a demonstrated causal streaming deployment.

8 Results

This section reports the executed results for the synthetic validation, LSS-PR-1.0 RD-domain evaluation, class-wise analysis, and IQ-domain ECA simulation.

Table I. Controlled synthetic smoke-test of detection methods: area under the receiver operating characteristic curve (AUC), probability of detection (Pd) at nominal false-alarm probability (Pfa), and runtime per frame. CUT, cell under test; CFAR, constant false-alarm rate; SVD, singular value decomposition; cICA, complex-valued independent component analysis; OR-CVCICA, online recursive complex-valued convolutive ICA. These simulation results are not the primary real-data evaluation.

Method	AUC	Pd@Pfa=0.1	Pd@Pfa=0.01	Runtime/frame (s)
Raw CUT energy	0.817	0.605	0.280	0.000014
CUT/reference ratio	0.817	0.605	0.280	0.000045
CA-CFAR-like	0.808	0.555	0.210	0.007024
SVD/whitening	0.993	0.990	0.895	0.000195
Batch cICA-lite	0.784	0.485	0.200	0.000649
OR-CVCICA	0.986	0.965	0.830	0.001845

The synthetic smoke-test in Table I and Fig. 2 checks pipeline behavior under controlled mixtures; it is not evidence for the primary real-data claim.

In the synthetic validation, OR-CVCICA achieved AUC 0.986 and Pd 0.830 at Pfa 0.01; SVD/whitening reached AUC 0.993. These numerical results characterize the synthetic task only.

Table II. Pooled detection performance on the LSS-PR-1.0 complex range–Doppler (RD)-domain benchmark. AUC, area under the receiver operating characteristic curve; Pd, probability of detection; Pfa, nominal false-alarm probability; Thr., score threshold; CUT, cell under test; CFAR, constant false-alarm rate. The CFAR-gated rank-score row represents the H0-calibrated, CFAR-gated online recursive complex-valued convolutive ICA (H0-CG-OR-CVCICA) detector.

Method	AUC	Runtime/frame (s)	Pd@Pfa=0.1	Thr@0.1	Pd@Pfa=0.01	Thr@0.01
Raw CUT energy	0.511	0.000013	0.254	124.472	0.101	170.206
CUT/reference ratio	0.511	0.000028	0.254	124.472	0.101	170.206
CA-CFAR-like	0.578	0.005819	0.287	84.259	0.133	174.397
SVD/whitening	0.340	0.000209	0.060	10209.665	0.011	24578.820
Batch cICA-lite	0.391	0.000705	0.113	12376.928	0.034	29696.375
OR-CVCICA-v1	0.402	0.002056	0.128	11470.612	0.047	23524.616
OR-CVCICA-v2 raw	0.439	0.004106	0.171	3.747	0.061	4.094
OR-CVCICA-v2 calibrated	0.561	0.004106	0.313	0.844	0.192	1.353
OR-CVCICA-v2 CFAR-gated rank score	0.719	0.004106	0.484	0.810	0.270	0.904

Table II reports the pooled CA-CFAR-assisted maximum-rank fusion as H0-CG-OR-CVCICA. The oriented component-derived score is normalized using H0 median and interquartile range (IQR), and the detector takes the maximum of the normalized empirical ranks of the calibrated component-derived score and the CA-CFAR-like score. This pooled evaluation uses offline score calibration and rank transformation.

The pooled evaluation produced AUC 0.719 for the reported proposed score and AUC 0.578 for the evaluated CA-CFAR-like baseline. At nominal Pfa values of 0.1 and 0.01, the proposed score yielded Pd 0.484 and 0.270, compared with 0.287 and 0.133 for CA-CFAR-like detection. These comparisons concern the implemented simplified baseline; they do not establish superiority over 2D CA-CFAR or OS-CFAR.

Table III. Target-class-specific area under the receiver operating characteristic curve (AUC) on LSS-PR-1.0 range–Doppler data. CUT, cell under test; CA-CFAR, cell-averaging constant false-alarm rate; SVD, singular value decomposition; cICA, complex-valued independent component analysis; OR-CVCICA, online recursive complex-valued convolutive ICA. Higher AUC indicates better class-specific discrimination.

Class	Raw CUT	CA-CFAR	SVD	Batch cICA	OR-v1	OR-v2 raw	OR-v2 cal.	H0-CG-OR-CVCICA
helicopter	0.365	0.575	0.147	0.161	0.169	0.192	0.808	0.898
passenger_ship	0.481	0.569	0.276	0.379	0.395	0.465	0.535	0.728
speedboat	0.641	0.580	0.480	0.512	0.518	0.542	0.458	0.584
uav	0.639	0.600	0.612	0.646	0.660	0.672	0.328	0.580

Class-specific evaluation is shown in Table III; pooled improvement does not imply improvement in every target class.

The proposed detector achieved class-wise AUC 0.898 for helicopters and 0.728 for passenger ships. For UAV, its AUC was 0.580, below 0.672 for OR-CVCICA-v2 raw. For speedboats, its AUC was 0.584, below 0.641 for raw CUT energy. Consequently, the qualitative UAV example in Fig. 9 is not quantitative evidence of superiority for the UAV class.

Table IV. Cancellation results in the separate simulated in-phase/quadrature (IQ)-domain experiment. ECA, extensive cancellation algorithm; P_in and P_out, input and residual output powers; dB, decibels. These simulated cancellation results are not direct comparisons with the real LSS-PR-1.0 range–Doppler benchmark.

Method	P_in	P_out	Cancellation (dB)	Runtime (s)	Dictionary size
batch ECA classical	4.076	0.006	28.017	0.610716	651
batch ECA reduced	4.076	0.012	25.444	0.048134	35
block ECA reduced	4.076	0.042	19.819	0.039568	35

These results are not directly compared with LSS-PR-1.0 because that release does not provide raw IQ reference/surveillance channels.

In the separate synthetic IQ experiment, classical batch ECA achieved 28.017 dB cancellation with dictionary size 651 and runtime 0.611 s; reduced ECA achieved 25.444 dB with dictionary size 35 and runtime 0.048 s; and block ECA achieved 19.819 dB and runtime 0.040 s. These comparisons are confined to the simulated IQ task.

8.1 Complexity sweep and runtime sensitivity

Twenty configurations were examined. At L_w = 2, K_V = 2, the reported AUC was 0.731, Pd at nominal Pfa 0.01 was 0.270 and runtime was 1.734 ms/frame. L_w = 2, K_V = 5 provided the same reported AUC of 0.731 at 1.651 ms/frame and was the selected operating point. These values belong to the complexity sweep, not to the independent repeated-validation mean.

Increasing the lag length did not improve detection performance. The mean AUC decreased from 0.731 at L_w = 2 to 0.717 at L_w = 9, while mean runtime increased from 1.678 ms/frame to 5.773 ms/frame. Throughput correspondingly decreased from approximately 596 frames/s to 173 frames/s. These results indicate that a short convolutive memory is sufficient for the current LSS-PR-1.0 RD-domain protocol, while longer lag embeddings mainly increase computational cost without improving detection accuracy.

Table V. Complexity-sweep results grouped by lag-embedding length (L_w) on LSS-PR-1.0 range–Doppler data. AUC, area under the receiver operating characteristic curve; Pd, probability of detection; Pfa, nominal false-alarm probability. Runtime and throughput summarize the tested whitening-update configurations for each L_w.

L_w	Mean AUC	AUC range	Pd@Pfa=0.01	Mean runtime (ms/frame)	Throughput (frames/s)
2	0.731	0.731-0.731	0.270	1.678	595.9
3	0.726	0.726-0.726	0.270	2.142	466.9
5	0.719	0.719-0.719	0.270	3.154	317.0
7	0.717	0.717-0.717	0.270	4.527	220.9
9	0.717	0.717-0.717	0.264	5.773	173.2

Table VI. Selected operating points from the LSS-PR-1.0 range–Doppler complexity sweep. L_w, lag-embedding length; K_V, whitening-update interval in frames; AUC, area under the receiver operating characteristic curve; Pd, probability of detection; Pfa, nominal false-alarm probability. "Recommended" and "Fastest" refer to the same tested configuration.

The three reported AUC values refer to different protocols: 0.719 is the pooled RD-domain evaluation (Table II); 0.731 is the short-lag complexity sweep operating point (Tables V–VI); and 0.714 ± 0.029 is the mean across 30 calibration/test partitions (Table VII). Neither the complexity sweep nor the pooled evaluation should be interpreted as an independent reproduction of the split-fitted result.

Configuration	L_w	K_V	AUC	Pd@Pfa=0.01	Runtime (ms/frame)	Comment
Best AUC	2	2	0.731	0.270	1.734	Highest AUC
Recommended	2	5	0.731	0.270	1.651	Balanced AUC-runtime
Fastest	2	5	0.731	0.270	1.651	Minimum runtime
Long lag example	9	5	0.717	0.265	5.683	Higher cost, no gain

8.2 Repeated file/group-level validation

Thirty repeated file/group-level calibration/test partitions were reported for 3423 observations. Each partition assigns 30 of the 43 H1 source files and 28 of the 40 H0 source-file/range-patch groups to calibration; 13 H1 files and 12 H0 patch groups remain for testing. The 30 manifest seeds are 2026–2055, with a nominal 0.3 test-group fraction. The group split does not duplicate a source-file/patch group across partitions, but the two H0 source recordings contribute different patches to calibration and testing, so this is not independent H0 source-file-held-out validation.

Table VII. Detection performance over 30 repeated file/group-level calibration/test splits on LSS-PR-1.0 complex range–Doppler (RD) data. AUC, area under the receiver operating characteristic curve; Pd, probability of detection; Pfa, false-alarm probability; CI, confidence interval; CFAR, constant false-alarm rate; OR-CVCICA, online recursive complex-valued convolutive independent component analysis. H0-CG-OR-CVCICA denotes the H0-calibrated, CFAR-gated OR-CVCICA detector. Reported ± values denote 95% CIs. Pfa = 0.01 is nominal; empirical Pfa is reported separately.

Method	AUC (mean ± 95% CI)	Pd at nominal Pfa = 0.01 (mean ± 95% CI)	Empirical Pfa	Runtime (ms/frame)
H0-CG-OR-CVCICA	0.714 ± 0.029	0.260 ± 0.031	0.015	1.320
OR-CVCICA-v2 calibrated	0.574 ± 0.038	0.199 ± 0.028	0.017	1.320
CA-CFAR-like	0.570 ± 0.015	0.132 ± 0.014	0.015	4.741

Method	AUC (mean ± 95% CI)	Pd at nominal Pfa = 0.01 (mean ± 95% CI)	Empirical Pfa	Runtime (ms/frame)
Raw CUT energy	0.508 ± 0.023	0.095 ± 0.010	0.012	0.009
CUT/reference ratio	0.508 ± 0.023	0.095 ± 0.010	0.012	0.022
OR-CVCICA-v2 raw	0.407 ± 0.032	0.046 ± 0.010	0.016	1.320
Batch cICA-lite	0.385 ± 0.034	0.035 ± 0.008	0.017	0.313
SVD/whitening	0.337 ± 0.032	0.014 ± 0.004	0.016	0.174

Across 30 file/group-level partitions, H0-CG-OR-CVCICA achieved a mean AUC of 0.714 ± 0.029, compared with 0.570 ± 0.015 for CA-CFAR-like detection, and mean Pd of 0.260 ± 0.031 versus 0.132 ± 0.014 at nominal Pfa 0.01. The empirical false-alarm probability of 0.015 exceeds the nominal level. The reported processing time is 1.320 ms per frame. These observations describe performance under the documented offline split protocol rather than end-to-end causal streaming.

8.3 Qualitative RD-domain illustration

Figure 9 shows a single selected UAV observation for visualization only. The target-like component index was chosen automatically by the OR-CVCICA-v2 component-scoring procedure, while the UAV frame itself was selected for illustration. The image depicts an ICA component–Doppler map rather than a reconstructed physical RD image; its localized response at Doppler bin 197 does not establish the reported detector's class-wise superiority or any quantitative improvement for UAV.

9 Figures and captions

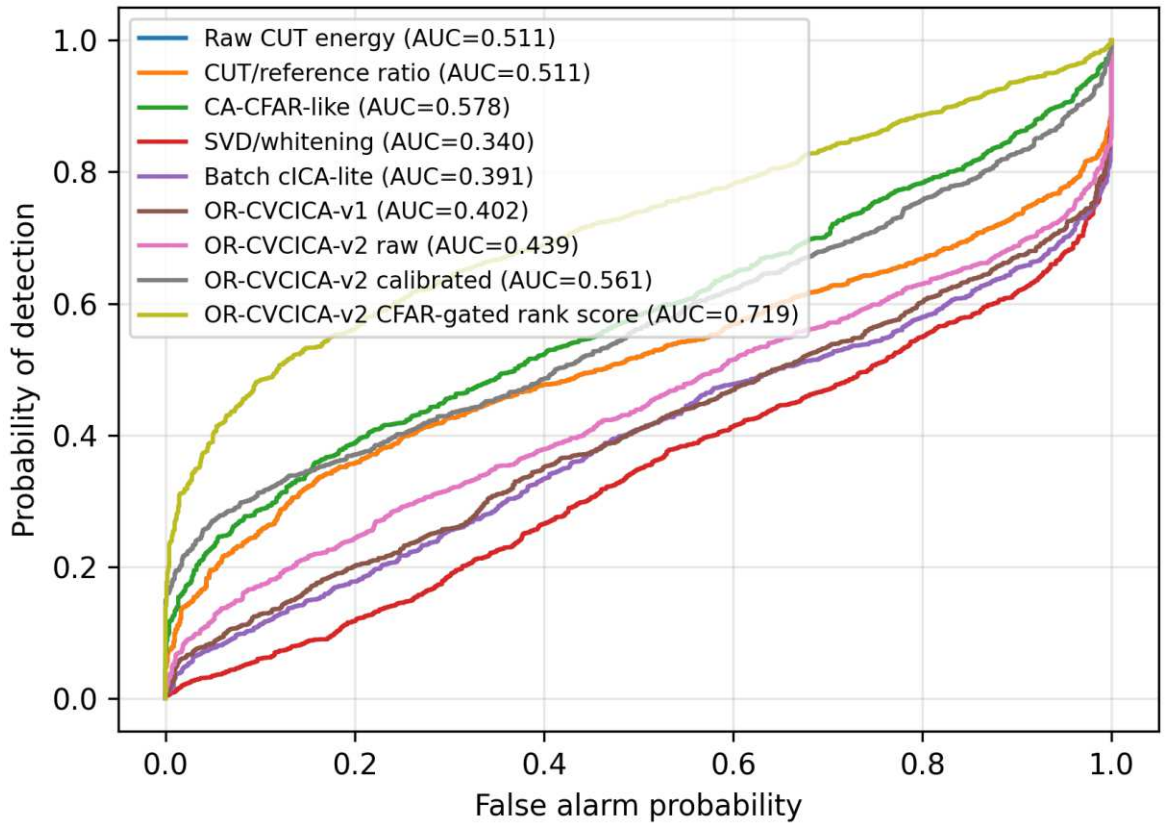


Fig. 1. Receiver operating characteristic (ROC) curves for the pooled evaluation on the LSS-PR-1.0 complex range–Doppler (RD)-domain benchmark. AUC denotes the area under the ROC curve. The H0-calibrated, CFAR-gated online recursive complex-valued convolutive ICA (H0-CG-OR-CVCICA) score has the highest overall AUC among the evaluated RD-domain methods; CFAR denotes constant false-alarm rate and ICA denotes independent component analysis.

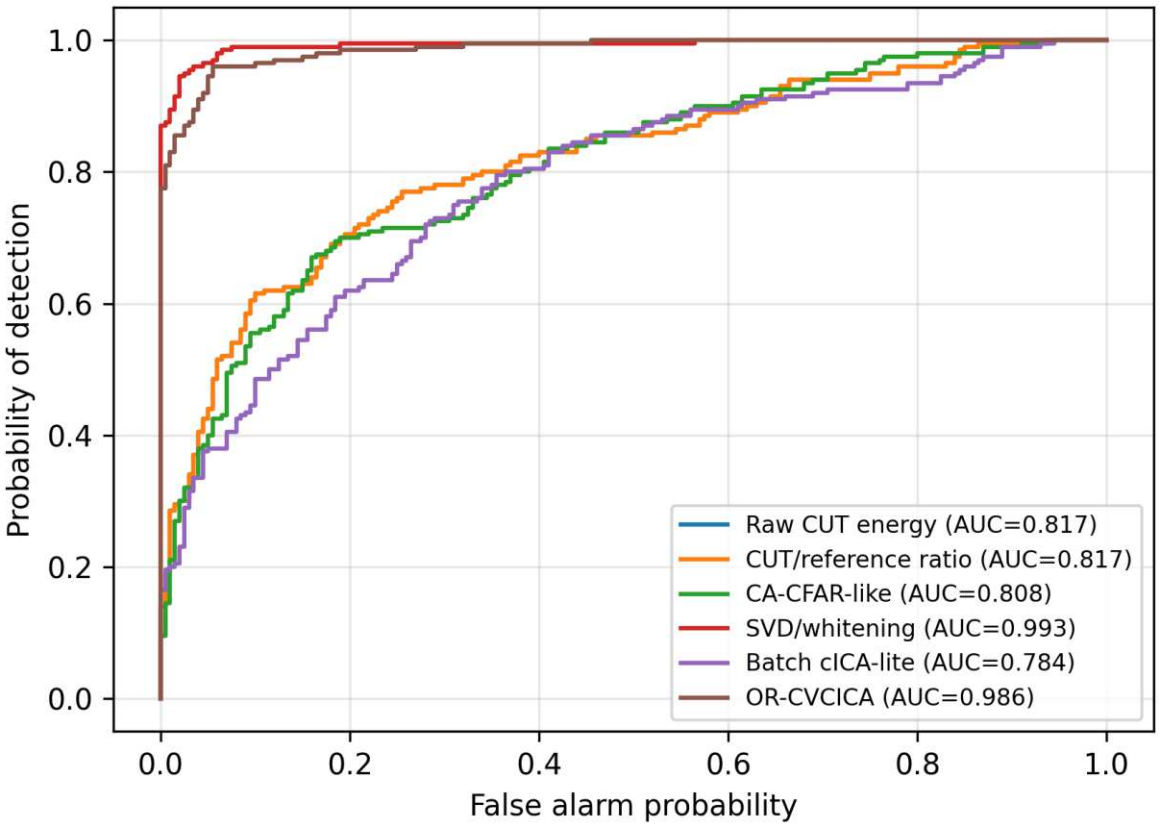


Fig. 2. Receiver operating characteristic (ROC) curves for the controlled synthetic smoke-test, with area under the curve (AUC) values shown in the legend. This simulation assesses the separation/detection pipeline under controlled mixture assumptions and is not the basis for the principal real-data performance claim. CUT denotes cell under test; CFAR, constant false-alarm rate; SVD, singular value decomposition; and ICA, independent component analysis.

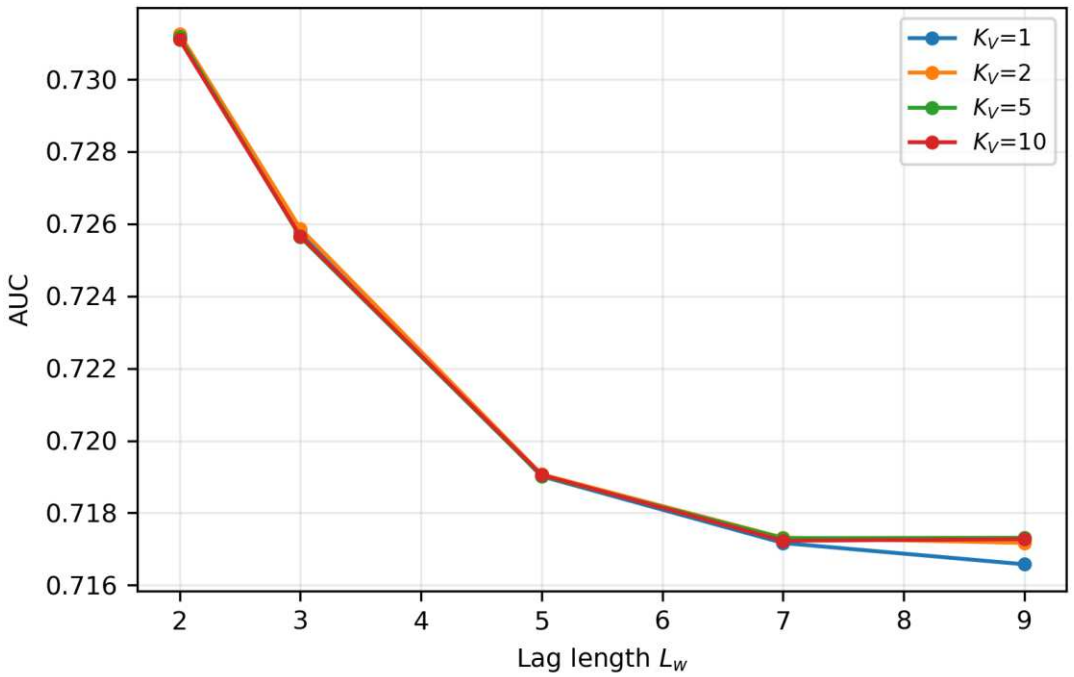


Fig. 3. Area under the receiver operating characteristic curve (AUC) versus lag-embedding length (L_w), with whitening-update interval K_V (in frames) as the curve parameter, in the LSS-PR-1.0 range–Doppler (RD)-domain complexity sweep. Higher AUC is observed at shorter L_w in the evaluated configurations.

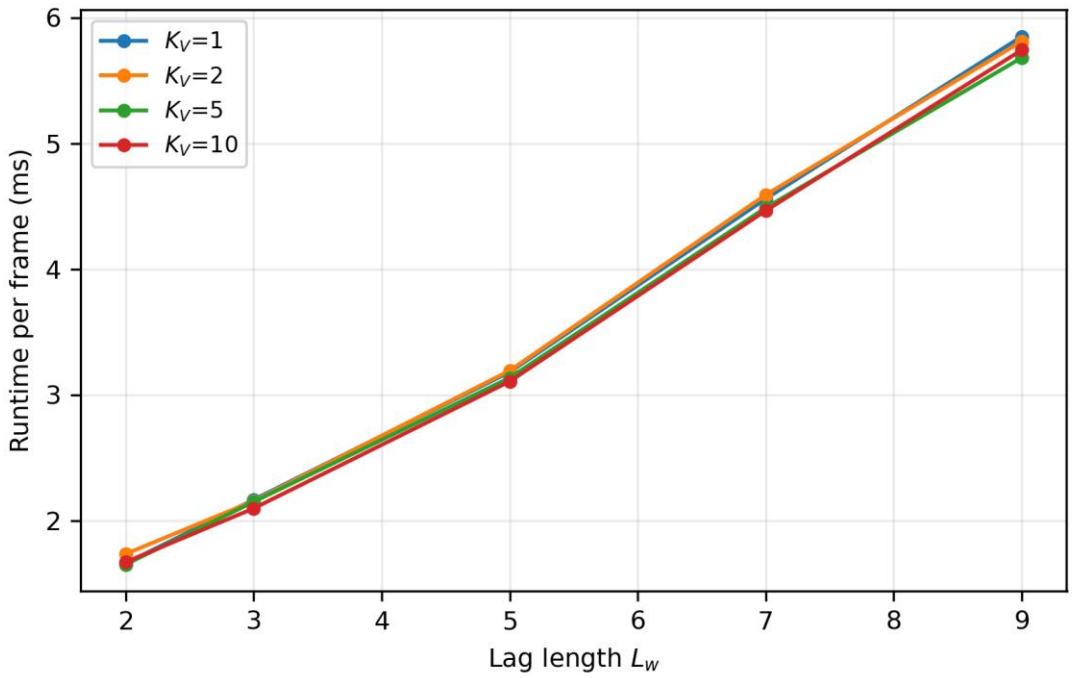


Fig. 4. Processing runtime per frame (milliseconds) versus lag-embedding length (L_w) for whitening-update intervals K_V (in frames) in the range–Doppler (RD)-domain complexity sweep. Runtime increases with L_w in the evaluated configurations.

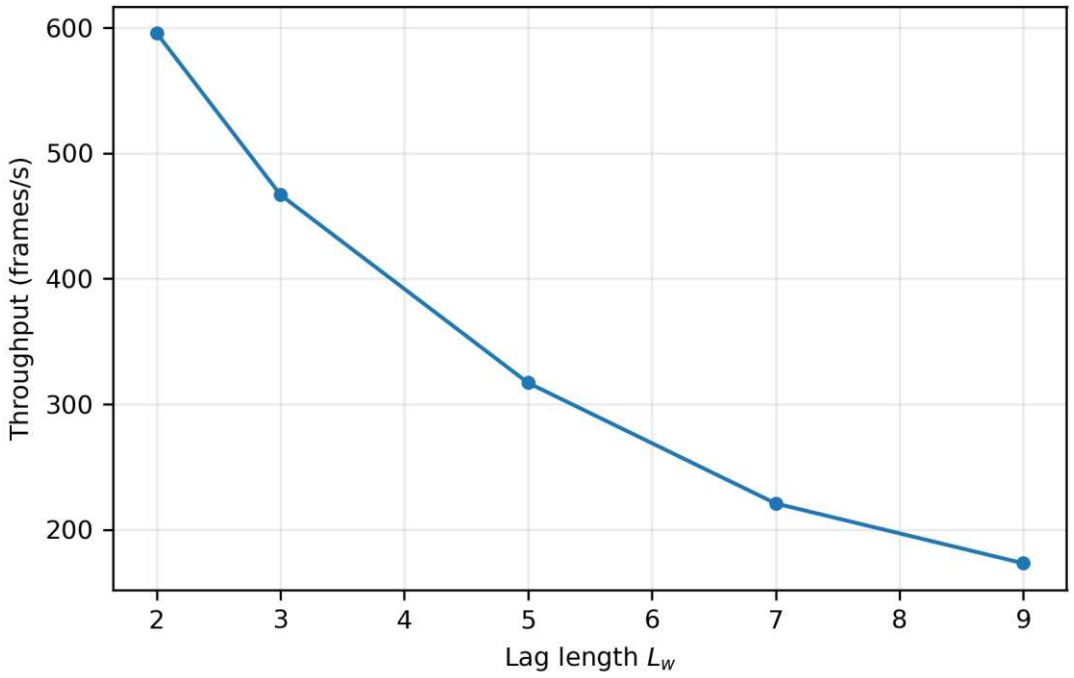


Fig. 5. Mean processing throughput (frames/s) versus lag-embedding length (L_w) across the evaluated whitening-update settings in the range–Doppler (RD)-domain complexity sweep. Throughput decreases as L_w increases in the tested implementation.

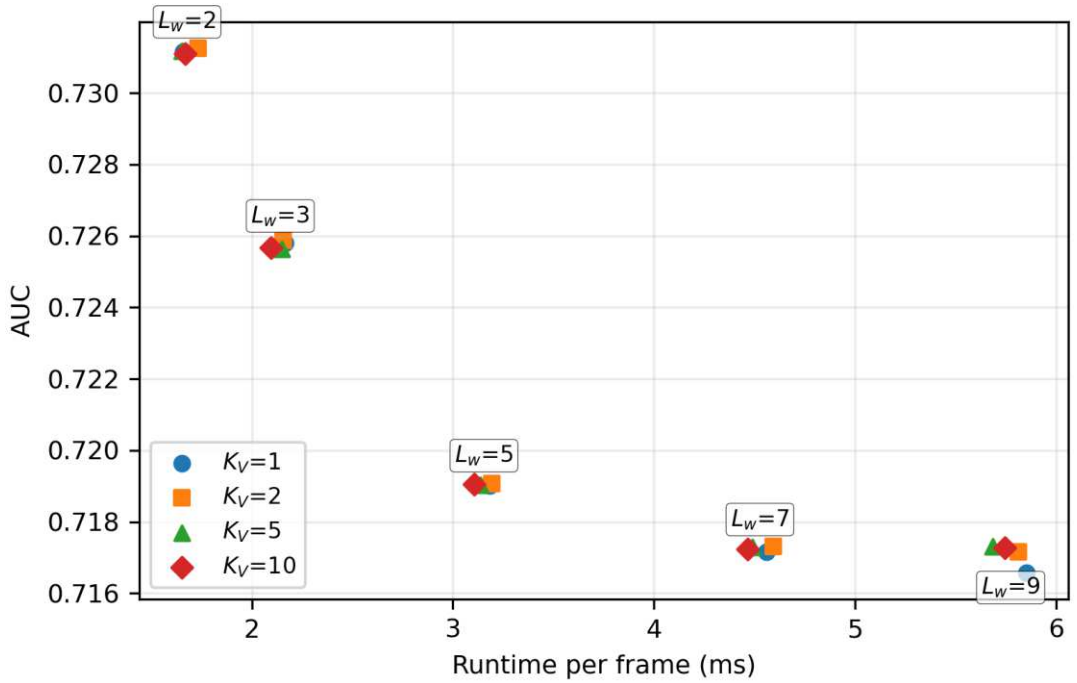


Fig. 6. Area under the receiver operating characteristic curve (AUC) versus runtime per frame (milliseconds) for the range–Doppler (RD)-domain complexity sweep. Labels indicate lag-embedding length L_w; marker types indicate whitening-update interval K_V (in frames). Shorter lags provide the strongest measured AUC–runtime trade-off.

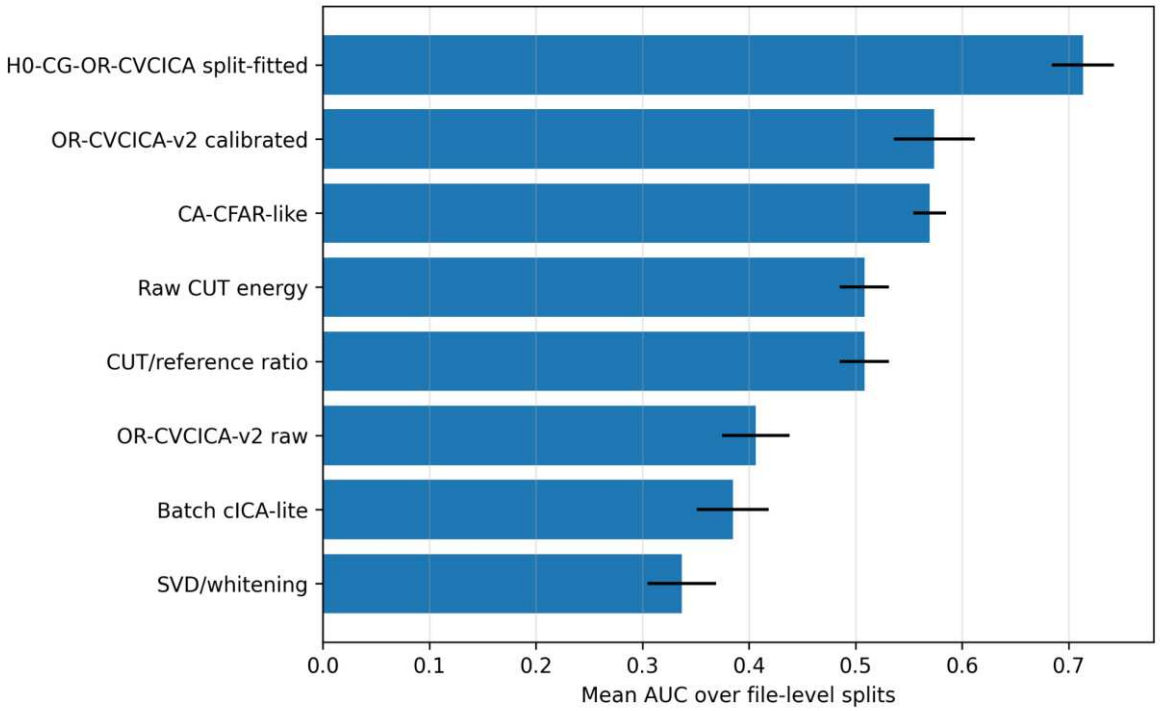


Fig. 7. Mean area under the receiver operating characteristic curve (AUC) across 30 repeated file/group-level calibration/test splits on LSS-PR-1.0 complex range–Doppler (RD) data. Horizontal error bars indicate 95% confidence intervals (CIs). H0-CG-OR-CVCICA denotes the H0-calibrated, constant false-alarm rate (CFAR)-gated online recursive complex-valued convolutive independent component analysis detector.

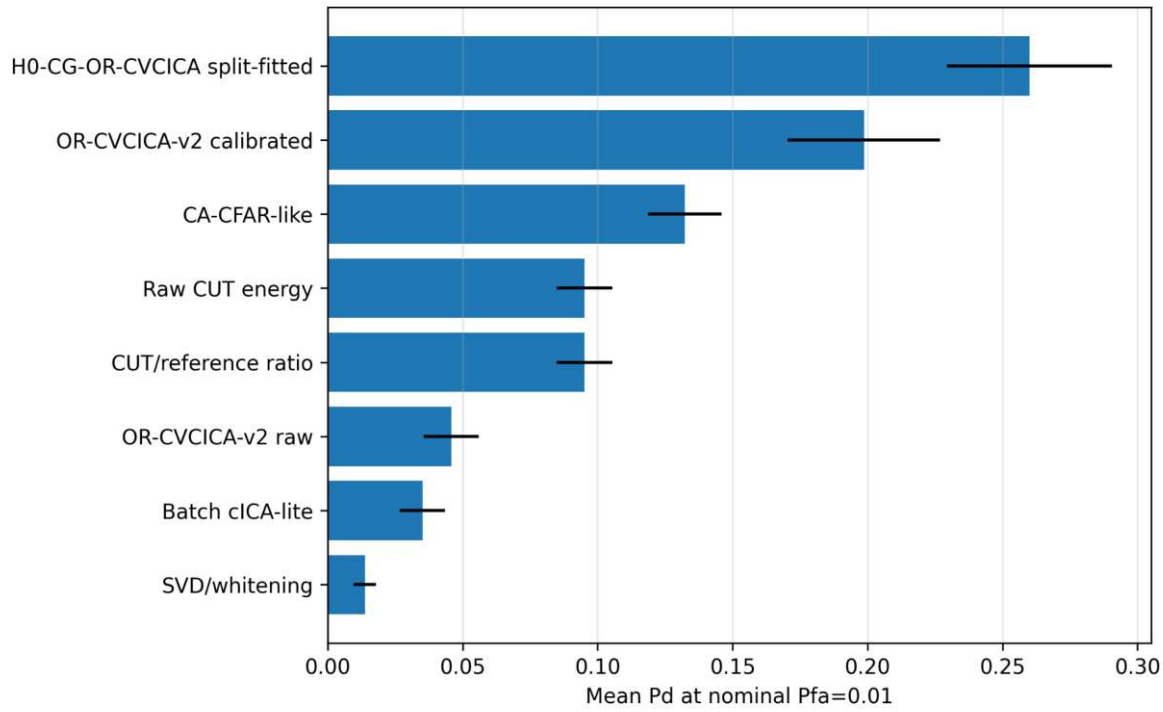


Fig. 8. Mean probability of detection (Pd) at nominal false-alarm probability Pfa = 0.01 across 30 repeated file/group-level calibration/test splits on LSS-PR-1.0 complex range–Doppler (RD) data. Horizontal error bars indicate 95% confidence intervals (CIs); the empirical Pfa may differ from the nominal value. CFAR denotes constant false-alarm rate; H0-CG-OR-CVCICA denotes the H0-calibrated, CFAR-gated online recursive complex-valued convolutive independent component analysis detector.

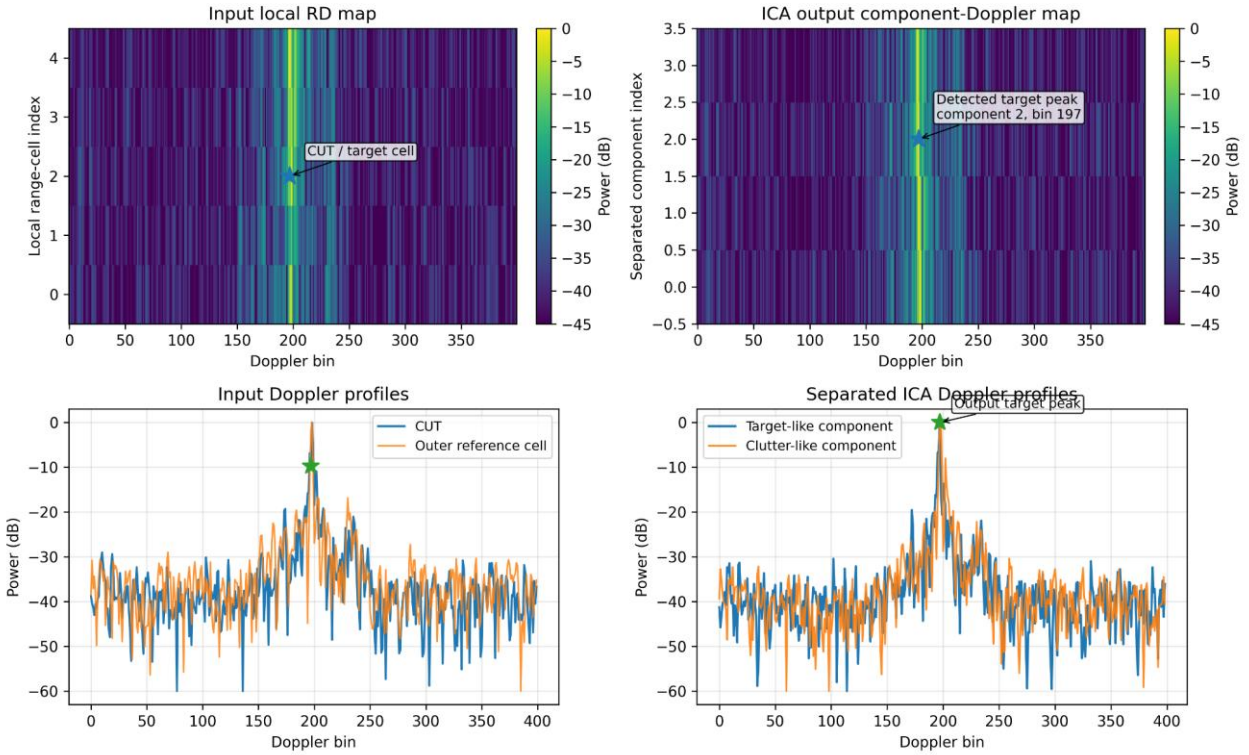


Fig. 9. Qualitative illustration of a selected unmanned aerial vehicle (UAV) frame from the LSS-PR-1.0 complex range–Doppler (RD) data, for visualization only. Upper left: input local RD map. Upper right: complex-valued independent component analysis (ICA) output shown as a component–Doppler map, not a physical RD reconstruction. The target-like component was selected automatically by the online recursive complex-valued convolutive ICA (OR-CVCICA-v2) component-scoring procedure, whereas the UAV frame was chosen for illustration. The star marks Doppler bin 197; lower panels show input cell-under-test (CUT)/reference profiles and separated component profiles. This single example does not establish class-wise detection superiority.

10 Discussion

The evaluated combination of component-derived scores and local clutter adaptation improves aggregate discrimination relative to the included one-dimensional CA-CFAR-like comparator. The pooled AUCs were 0.719 versus 0.578. In 30 file/group-level partitions, mean AUCs were 0.714 ± 0.029 versus 0.570 ± 0.015, and mean Pd at nominal Pfa 0.01 was 0.260 ± 0.031 versus 0.132 ± 0.014. The empirical Pfa of 0.015 for the proposed score exceeds its nominal level and does not prove strict CFAR control.

Score orientation requires representative labeled H0 and H1 calibration observations. Target-absent data supply the IQR-based robust normalization reference and threshold distribution, but they do not independently determine the target-present direction of the score. Thus, the approach is not purely unsupervised and cannot be described as H0-only in every calibration step.

Aggregate benefit was not uniform across classes. Helicopter and passenger-ship AUCs were 0.898 and 0.728 for the proposed method, while for UAV the proposed AUC was 0.580 against 0.672 for the raw OR-CVCICA-v2 score. For speedboats, raw CUT energy achieved 0.641 against 0.584 for the proposed score. This motivates class-dependent fusion analysis rather than a claim of universal detection superiority.

Figure 9 is a qualitative, non-quantitative visualization of one selected UAV frame. Component selection was automatic within that observation, but selection of the record/frame was for visualization only. The illustrated ICA component–Doppler output is not a physical range–Doppler reconstruction and does not resolve the UAV class-wise AUC deficit.

Only two independent H0 sea-clutter recordings contribute the 1040 negative observations; repeated frames and patches do not create additional independent recording conditions. Group-level splitting allows different H0 patches from the same recordings in calibration and testing, and acquisition matching between H1 and H0 files has not been established. Generalization to previously unseen clutter recordings therefore remains uncertain.

The complexity sweep AUC of 0.731 at $L_w = 2$ is distinct from the pooled AUC 0.719 and repeated split mean AUC 0.714. The selected sweep configuration achieved 1.651 ms/frame, whereas repeated validation reports 1.320 ms/frame; these are protocol-specific measurements. They indicate computationally moderate processing in the tested Python implementation, not demonstrated end-to-end real-time performance.

The recursive covariance, whitening and demixing stages use the current observation and retained state, whereas the complete detection experiments were carried out offline. Pooled score orientation and rank normalization are dataset-level operations, while the repeated file/group-level protocol reports calibration/test partitions. The method employs IQR-based normalization and maximum-rank fusion, without a separate scalar gate weight. The available split manifests document the repeated evaluation but do not establish operational performance under end-to-end causal streaming. Future investigation should include independent clutter recordings, stronger CFAR baselines, fully specified execution environments, and hardware-aware streaming experiments. The real-data conclusions are restricted to processed RD observations; IQ-domain ECA comparisons were performed in simulation.

11 Conclusion

This work investigated H0-CG-OR-CVCICA for clutter-aware weak-target detection in processed complex passive-radar RD observations. In the pooled LSS-PR-1.0 evaluation, the reported AUC was 0.719 and Pd at nominal Pfa 0.01 was 0.270, compared with 0.578 and 0.133 for the evaluated CA-CFAR-like baseline. Across 30 repeated file/group-level partitions, mean AUC was 0.714 ± 0.029 and mean Pd at nominal Pfa 0.01 was 0.260 ± 0.031, with empirical Pfa 0.015. Processing time was reported as 1.320 ms/frame in the tested Python setting; this does not demonstrate a fully causal streaming system. Real-data conclusions apply only to complex RD observations. Separate IQ simulations addressed ECA-type cancellation and are not direct LSS-PR-1.0 comparisons. The detected aggregate benefit was class-dependent, relied on only two H0 source recordings and was assessed against a limited baseline set. Future work should include independent clutter recordings, robust gate/score reproducibility, stronger CFAR baselines and end-to-end streaming tests.

The illustrative UAV component–Doppler map in Fig. 9 provides a visualization of the separation stage only and does not establish class-wise detection superiority.

Appendix A. Reproducibility and implementation details

A.1 Data preparation and recursive separation

All real observations are complex RD samples. There are 43 target-present source files (2383 frames) and two H0 sea-clutter source files (1040 range patches over 52 frames). Each frame has $D = 400$ Doppler bins; $M = 2$ constructed channels consist of the central CUT and the average of the two outer cells of a five-cell patch. H0 patch start indices are 0, 5, 10, …, 95 over 101 range cells. The input channels are separately centered and normalized to unit average power. A length-L_w lag embedding produces $P = M L_w$ components and $T = D - L_w + 1$ embedded Doppler samples per frame.

Within the defined observation order, covariance and demixing retain a recursive state: $R_k = \lambda_R R_{(k-1)} + (1-\lambda_R)U_k U_k^H/T$; whitening periodically uses the inverse Hermitian square root of $R_k + \varepsilon I_P$; and natural-gradient demixing updates W_k with $\varphi(Y) = Y/(|Y|+\rho)$, followed by symmetric decorrelation. These are frame-sequential state operations, not end-to-end verified real-time radar processing.

A.2 Processing pipeline (implementation-oriented pseudocode)

1 Construct RD mixture: for H1 frames use the central local range cell as CUT and mean of the two outer cells as REF; for H0 iterate 20 disjoint five-cell range patches per frame.

2 Independently center and normalize the two complex channels; form X_k with shape 2 × 400.

3 Lag-embed each current frame X_k over L_w adjacent Doppler bins, obtaining U_k with P = 2L_w rows and T = 400−L_w+1 columns.

4 Update $R_k \leftarrow \lambda_R R_{(k-1)} + (1-\lambda_R)U_k U_k^H/T$; enforce Hermitian symmetry and diagonal regularization.

5 At the scheduled K_V refresh, compute adaptive whitening V_k; compute $Z_k \leftarrow V_k U_k$.

6 Calculate $Y_k \leftarrow W_k Z_k$ and $\varphi(Y_k) = Y_k/(|Y_k|+\rho)$; apply natural-gradient update with μ_k and symmetric decorrelation to obtain the next W_k.

7 Extract the component descriptors and automatically select the target-like index maximizing the specified weighted component score; produce raw $s_{or,k}$.

8 Independently compute the one-dimensional Doppler CA-CFAR-like frame score $s_{cfar,k}$ with 24 training and four guard bins per side.

9 Fit a_cal using labeled H0/H1 calibration observations (reverse the raw score when calibration AUC < 0.5). For the oriented H0 scores, fit the median and interquartile range (IQR).

10 For each observation, apply score orientation and median/IQR normalization, then compute $r_{or,k}$ and $r_{cfar,k}$ by normalized empirical ranks; form the fused statistic $T_k \leftarrow \max(r_{or,k}, r_{cfar,k})$. For the pooled v2 experiment these ranks are calculated from the pooled score vectors; reported repeated validation uses calibration/test partitions.

11 Estimate $\eta(Pfa)$ as the empirical $(1-Pfa)$ quantile of H0 fused scores and decide $\delta_k \leftarrow 1\{T_k > \eta(Pfa)\}$; report AUC, Pd and empirical Pfa on the designated evaluation partition.

The separation and scoring implementation uses H0 median/IQR normalization and maximum-rank fusion, without a separate gate-weight parameter. The repeated validation is documented through the reported sample/split manifests and aggregate results. The reported partition-level statistics should be distinguished from an independently benchmarked end-to-end streaming detector.

A.3 Hyperparameters and selection

Table A1. Hyperparameters and implementation settings for the evaluated processing configurations. Pooled and repeated-split settings are identified separately where their configurations differ.

Parameter	Reported value / interpretation	Status / selection
M; D	2 channels; 400 Doppler bins	RD input construction
L_w; K_V	2; 5 for repeated validation; 5; 1 in pooled v2 code	Separate experiment configurations
Sweep values	L_w = 2,3,5,7,9; K_V = 1,2,5,10	20 tested settings
λ_R; μ	0.98; 0.02	v2 implementation defaults
ρ; ε_R	10^{-6}; 10^{-6}	v2 nonlinearity and covariance regularization
ε_{score}	10^{-12}	Baseline denominators; score-calibration numeric stabilization
Components	P = ML_w; four at L_w = 2	Demixing state dimension
Target component	Automatic argmax weighted component score	Automatic component selection in v2 implementation
Component weights	Contrast .35; compactness .20; decorrelation .25; non-Gaussianity .10; temporal consistency .10	v2 component scoring settings
Fusion operation	$T_k = \max(r_{or,k}, r_{cfar,k})$; no g_k parameter	Maximum-rank fusion used in pooled v2 scoring
H0 normalization	Median / interquartile range (IQR) of oriented H0 scores	H0 calibration statistics from v2 implementation
Orientation	a_cal = −1 if labeled calibration AUC < 0.5; otherwise +1	ScoreCalibrator.fit() rule
CA-CFAR-like	24 training + 4 guard bins per Doppler side	One-dimensional baseline implementation
Nominal Pfa	0.1 and 0.01	Empirical H0 quantile thresholds
Validation repetitions	30; seeds 2026–2055	Reported final split manifest

A.4 Repeated calibration/test partition protocol

There are 83 grouping units: 43 H1 source-file groups and 40 H0 source-file/range-patch groups. In each of the 30 splits (seeds 2026 to 2055), calibration uses 30 H1 files and 28 H0 patch groups, whereas testing uses 13 H1 files and 12 H0 patch groups. The nominal test-group fraction is 0.3. Calibration has 2162–2632 observations and testing 791–1261 depending on group size. Both

original H0 source files contribute different patch groups to calibration and testing; source-file-independent H0 generalization was not evaluated.

The sample and split manifests document the reported partition membership, sample counts, and corresponding aggregate detection results. Performance is interpreted within this repeated offline file/group-level calibration/test protocol; it is not presented as independent source-recording validation for H0 or as an end-to-end streaming benchmark.

A.5 Computing environment and timing

The research implementation uses Python 3 with NumPy (≥1.23), SciPy (≥1.9), h5py (≥3.8), pandas (≥1.5), Matplotlib (≥3.6), and scikit-learn (≥1.2). The pooled timing procedure uses time.perf_counter() to time sequential score generation and reports elapsed time divided by the observation count. The reported per-frame timing excludes the full data-loading, offline calibration, rank-normalization, and threshold-fitting workflow. The 30 repeated calibration/test partitions represent statistical validation splits, not 30 separately timed hardware benchmarks.

A.6 Data and code availability

Data availability. The real-data RD observations derive from the public LSS-PR-1.0 dataset, Science Data Bank, https://doi.org/10.57760/sciencedb.14468 [12]. Only the processed complex RD release was used for real-data detection; separate synthetic IQ recordings were used only for ECA-type cancellation comparisons.

Code availability. Python implementation files for the recursive ICA separator, RD baselines, IQR score calibration, maximum-rank fusion, simulation, and visualization are available from the corresponding author upon reasonable request. The sample/split manifests and aggregate validation summaries document the repeated partition-level evaluation. Correspondence regarding research software and supporting execution records should be directed to Hamze Ghahramani at the email addresses on the title page.

Statements and declarations

Funding. No funding was received for this study.

Competing interests. The author declares no competing interests.

Author contributions. Hamze Ghahramani is the sole author and is responsible for the study conception and design, implementation, experiments, analysis, interpretation, manuscript preparation, review, and approval of the final manuscript.

Declaration of AI-assisted work. ChatGPT (OpenAI) was used to assist with coding, experimental design, manuscript preparation, and English-language editing and review. The sole author carefully and comprehensively reviewed the experiments and manuscript and accepts full responsibility for the content, accuracy, and integrity of the work.

Ethics approval. Not applicable. This study used publicly available radar observations and simulated signals and did not involve human participants or animals.

Consent to participate and consent for publication. Not applicable; no human participants or identifiable personal data were involved.

Data availability and code availability. See Appendix A.6 for dataset DOI, processing scope and the status of reproducibility code.

References

[1] Colone F, O'Hagan DW, Lombardo P, Baker CJ (2009) A multistage processing algorithm for disturbance removal and target detection in passive bistatic radar. IEEE Trans Aerosp Electron Syst 45(2):698–722. https://doi.org/10.1109/TAES.2009.5089551

[2] Colone F, Palmarini C, Martelli T, Tilli E (2016) Sliding extensive cancellation algorithm for disturbance removal in passive radar. IEEE Trans Aerosp Electron Syst 52(3):1309–1326. https://doi.org/10.1109/TAES.2016.150477

[3] Garry JL, Baker CJ, Smith GE (2017) Evaluation of direct signal suppression for passive radar. IEEE Trans Geosci Remote Sens 55(7):3786–3799. https://doi.org/10.1109/TGRS.2017.2680321

[4] Sui J, Wang J, Zuo L, Gao J (2023) Multistage least squares algorithms for clutter suppression in airborne passive radar based on subband operation. IEEE Trans Aerosp Electron Syst 59(2):1893–1909. https://doi.org/10.1109/TAES.2022.3207126

[5] Blasone GP, Colone F, Lombardo P, Wojaczek P, Cristallini D (2021) Passive radar STAP detection and DoA estimation under antenna calibration errors. IEEE Trans Aerosp Electron Syst 57(5):2725–2742. https://doi.org/10.1109/TAES.2021.3061803

[6] Nikaein H, Ghaderi Zefreh R, Gazor S (2023) Target detection and interference cancellation in passive radar sensors via group-sparse regression. IEEE Sens J 23(18):21484–21492. https://doi.org/10.1109/JSEN.2023.3300887

[7] Nikaein H, Sheikhi A, Gazor S (2021) Target detection in passive radar sensors using least angle regression. IEEE Sens J 21(4):4533–4542. https://doi.org/10.1109/JSEN.2020.3035630

[8] Xie D, Yi J, Wan X, Jiang H (2023) Masking effect mitigation for FM-based passive radar via nonlinear sparse recovery. IEEE Trans Aerosp Electron Syst 59(6):8246–8262. https://doi.org/10.1109/TAES.2023.3301459

[9] Ghahramani H, Parhizgar N, Abbasi Arand B, Barari M (2021) Convolutive complex-valued independent component analysis for nonlinear radar signal processing and maritime weak target detection. Math Probl Eng 2021:6191303. https://doi.org/10.1155/2021/6191303

[10] Hyvärinen A, Oja E (2000) Independent component analysis: algorithms and applications. Neural Netw 13(4–5):411–430. https://doi.org/10.1016/S0893-6080(00)00026-5

[11] Chen X, Rao G, Guan J, Wang J, Wang H, Zhang C, Yi J, Wan X, Rao Y (2025) Passive radar low slow small detection dataset (LSS-PR-1.0) and multi-domain feature extraction and analysis methods. J Radars 14(2):249–268. https://doi.org/10.12000/JR24145

[12] Chen X, Guan J, Yi J (2026) LSS-PR-1.0: Passive Radar Sea and Air Targets Detection Dataset. Science Data Bank. https://doi.org/10.57760/sciencedb.14468

[13] Cui K, Wang C, Zhou F, Liu C, Gao Y, Feng W (2024) Fast algorithm of passive bistatic radar detection based on batches processing of sparse representation and recovery. Remote Sens 16(13):2294. https://doi.org/10.3390/rs16132294

[14] Javidan MH, Zaimbashi A, Liu J (2020) Target detection in passive radar under noisy reference channel: a new threshold-setting strategy. IEEE Trans Aerosp Electron Syst 56(6):4711–4722. https://doi.org/10.1109/TAES.2020.2999998

[15] Zuo L, Wang J, Zhao T, Cheng Z (2021) A joint low-rank and sparse method for reference signal purification in DTMB-based passive bistatic radar. Sensors 21(11):3607. https://doi.org/10.3390/s21113607

[16] Liu Y, Yi J, Wan X, Zhang X, Ke H (2020) Time-varying clutter suppression in CP-OFDM based passive radar for slowly moving targets detection. IEEE Sens J 20(16):9079–9090. https://doi.org/10.1109/JSEN.2020.2986717

[17] Nikaein H, Sheikhi A, Gazor S (2021) Multitarget detection in passive MIMO radar using block sparse recovery. IEEE Access 9:121206–121216. https://doi.org/10.1109/ACCESS.2021.3108195

[18] Zhou Q, Yuan Y, Venturino L, Yi W (2024) Direct target localization for distributed passive radars with direct-path interference suppression. IEEE Trans Signal Process 72:3611–3625. https://doi.org/10.1109/TSP.2024.3414423

[19] Sun Q, Shan T, Feng Y, He S, Bai X, Zhang H (2024) A frequency domain clutter suppression approach for passive radar. IEEE Sens J 24(9):14916–14929. https://doi.org/10.1109/JSEN.2024.3382103

[20] Blázquez-García R, Cristallini D, Ummenhofer M, Seidel V, Heckenbach J, O'Hagan D (2024) Capabilities and challenges of passive radar systems based on broadband low-Earth orbit communication satellites. IET Radar Sonar Navig 18(1):78–92. https://doi.org/10.1049/rsn2.12446

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformationH0CGORCVCICA.docx