

Physics-Informed Dual-Modality Monitoring for GNSS Spoofing Detection

Liupeng Chen

National University of Defense Technology

Gang Ou

National University of Defense Technology

Zhen Huang

National University of Defense Technology

Zheng Wei

National University of Defense Technology

Weihua Mou

whmou@nudt.edu.cn

National University of Defense Technology

Research Article

Keywords: GNSS spoofing detection, Signal quality monitoring, Physics-informed learning

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11233778/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Physics-Informed Dual-Modality Monitoring for GNSS Spoofing Detection

Liupeng Chen^{1,2}, Gang Ou^{1,2}, Zhen Huang^{1,2}, Zheng Wei^{1,2} and Weihua Mou^{1,2,*}

¹College of Electronic Science and Engineering, National University of Defense Technology, Changsha 410073, China

²National Key Laboratory of Navigation and Spatiotemporal Technology, Changsha 410073, China

*Corresponding author: Weihua Mou (whmou@nudt.edu.cn)

Abstract

Global Navigation Satellite Systems (GNSS) underpin positioning, navigation, and timing (PNT), but open civil signals remain exposed to jamming and spoofing. When neither the radio-frequency configuration nor the attack is known in advance, the receiver must still identify the spoofed satellite channels before those measurements can be mitigated. Physics-informed dual-modality cross-attention networks (PI-DMCANs) are introduced for this channel-wise task. A short tracking-feature sequence and a local cross-ambiguity-function image are fused by causal cross-attention, and a residual bank—code-carrier consistency, signal-quality monitoring (SQM), peak walk, a secondary peak, and clock drift—sets the weight of each stream. On TEXBAT, OAKBAT, and the multi-constellation dataset MuCoGS, PI-DMCANs were compared with five baselines. Within-scenario accuracy at a Neyman–Pearson threshold was 0.965. Under within-dataset cross-scenario evaluation and cross-dataset evaluation, accuracy was 0.928 and 0.932 respectively, about 3.5 and 5.1 percentage points above the strongest baseline. Ablations of modality fusion and of the gate confirm that dual-modality residual gating is most helpful under hold-out. The resulting channel flag marks which satellite a GNSS receiver should de-weight or exclude, and thereby supplies a basis for interference mitigation.

Keywords: GNSS spoofing detection; Signal quality monitoring; Physics-informed learning;

1 Introduction

The coexistence of GPS, GLONASS, Galileo, and BeiDou with regional augmentation improves geometry, redundancy, and availability relative to single-constellation operation, and it supports PNT applications with strict integrity requirements, including automated driving, disaster response, and critical infrastructure. The same multi-constellation measurements have driven PPP/PPP-RTK and bias calibration (Cui et al. 2023; Wu et al. 2024). Redundancy alone does not remove radio-frequency threats. Civil signals are weak and openly specified, so receivers in urban canyons, ports, and near critical facilities remain exposed to unintentional interference, jamming, and spoofing (Psiaki and Humphreys 2016; Radoš et al. 2024). Spoofing is the more damaging of these threats because it does not simply deny a fix: fabricated code phase, carrier, and navigation data pull the tracking loops onto a prescribed trajectory (Gao et al. 2020), and integrity can be lost while the position still appears nominal. User-side risk control, including de-weighting of anomalous observations, is therefore required (Wang et al. 2025). On a single-antenna receiver without an inertial unit, array direction finding, dual-antenna carrier-phase differences, and INS-aided monitoring are unavailable (Liu et al. 2023; Zhang et al. 2024; Jin et al. 2025). Interference mitigation and integrity alerting must then be carried out channel by channel in the post-correlation tracking domain. A chi-square test on raw intermediate-frequency samples can report that an anomaly is present, but it does not name the satellite channel that should be de-weighted (Bhuiyan et al. 2026). Carrier-to-noise density ratio (C/N_0), multi-correlator SQM, code-carrier

consistency, clock dynamics, and the local complex cross-ambiguity function (CCAF) are the observables used to decide whether each visible satellite is nominal (Wesson et al. 2018; Islam et al. 2024).

The tracking domain already provides these observables. The difficulty is that they are heterogeneous, are not synchronous, and fail in turn as the attack or the constellation changes. An overpowered takeover first perturbs automatic gain control and C/N_0 . A matched-power pull-off then holds the power advantage near 1dB for a long interval, so correlation-peak distortion appears late and remains weak. A jam-then-spoof sequence is different again: after the tracking outage it reappears already false-locked (Humphreys et al. 2012; Wesson et al. 2018). Concatenating the observables and training a support vector machine (SVM) or a shallow network can still be accurate when that scenario has been seen. Once an entire record, the front end, or the constellation is held out, however, the decision surface often memorizes front-end fingerprints and scenario labels rather than transferable physical structure (Semanjski et al. 2019; Chen et al. 2022; Iqbal et al. 2024a). The difficulty is greater for a multi-constellation receiver, because clean medians already differ between BPSK and BOC families, and false-alarm calibration across corpora is stricter than a split inside one dataset.

Recent work has advanced along three lines that are seldom combined in one monitor. SQM has been treated as a random process and filtered or linearly combined: a moving variance emphasizes the takeover transient (Sun et al. 2018a); multi-metric combinations and quadrature-branch monitoring improve robustness (Sun et al. 2018b; 2021); and indirect Kalman filtering (IKF-SQM) sets a threshold from the steady-state variance of a single SQM statistic (Jin et al. 2026b). Aggregated correlation-residue likelihood analysis (A-CoRLiAn) provides a lightweight decision on TEXBAT with a small number of extra correlators (Ji et al. 2024). A second line treats correlator lag versus time as an image. Jin et al. define a spoofing signal-to-noise ratio and show that a fixed linear mix of SQM metrics cannot attain it; SpoofConvNet therefore replaces that mix with a convolutional encoder on the correlator grid (Jin et al. 2026a). Deep networks on CAF patches have been used to count and classify peaks (Borhani-Darian et al. 2024). Sparse decomposition can separate authentic and counterfeit components inside the tracking loop and can support position recovery. A third line writes a physical relation into the loss. Code-carrier consistency physics-constrained neural networks (CCC-PCNNs) keep held-out TEXBAT and OAKBAT accuracy close to the matched-scenario figure, yet the penalty is one code-carrier relation, the input is a one-dimensional sequence, and multi-constellation jam-then-spoof scenarios are outside that study (Song et al. 2026).

Three gaps follow from this survey. First, no single physical quantity covers every attack family. In the middle of a matched-power interval, both code-carrier consistency and the peak shape can lie close to the clean distribution, whereas overpowered capture and common-mode loss of lock are conspicuous. A weak antenna or multipath, by contrast, raises SQM on a clean channel. Penalizing every sample so that its residual must resemble spoofing therefore inflates the false-alarm rate. Second, the tracking sequence and the CCAF image are not homologous. A local magnitude and quadrature (Mag/Q) heatmap on the code-delay axis and a short tracking sequence differ by about an order of magnitude in time scale, and their missingness patterns differ, so they cannot be treated as two copies of the same stream or simply concatenated. Mapping monitoring quantities into images has precedent in wide-area C/N_0 interference localization (Chen Q et al. 2025), but a transferable fusion of the short sequence with the CCAF image is still lacking for a single receiver. Third, public datasets and constellation families do not share a radio-frequency configuration. Front ends, scenario designs, and signal structures all differ, so the C/N_0 floor and the clean Ratio median both drift. An unseen-scenario split inside one dataset therefore does not guarantee that the false-alarm rate stays calibrated on another configuration. Following the use of a filter steady-state variance for thresholding in IKF-SQM (Jin et al. 2026b), the output is evaluated under a

Neyman–Pearson (NP) rule: within-scenario and cross-scenario tables share the empirical 95th percentile of clean training scores with the five baselines, so that the operating point remains comparable.

The literature already contains many monitors for GNSS integrity. It does not yet contain one channel-wise monitor that joins multi-residual conditional gating, a one-sided score–residual consistency term, and causal cross-attention between the two modalities. The method proposed here supplies that monitor. For each visible satellite it produces a normal or anomalous decision at a controlled false-alarm rate. The flag is then a basis for interference mitigation: the affected channel can be de-weighted, isolated, or excluded before it enters the navigation solution.

The main contributions of this paper are summarized as follows:

- **Dual-modality cross-attention monitoring.** Useful evidence does not always lie in the same modality. A short tracking sequence provides the tracking features, and a local Mag/Q heatmap provides the CCAF image. Tracking tokens are the queries and CCAF tokens are the keys and values. Causal cross-attention then fuses the two representations into a channel-wise spoofing score.
- **Gating by a physical residual bank.** Adding the two modalities with equal weight lets the less informative stream misleads the decision. Code–carrier divergence, the SQM level and its short-term jitter, the code-chip offset of the main peak from zero delay, a far-sidelobe secondary peak, and the relative clock-drift offset are therefore collected in a residual bank that sets the two stream weights. A learnable soft prior is added, so the more informative modality receives a larger weight.
- **Cross-constellation generalization.** Code rates and signal structures differ across constellations, so clean peak shapes, SQM medians, and anomaly-score floors cannot be shared. Code–carrier Doppler residuals are first converted to range with the carrier and code reference frequencies of each constellation. Ratio medians and residual percentiles are then estimated on clean training windows of that constellation family. The residual references therefore stay constellation-specific, while the detector itself is shared.

The remainder of the paper is organized as follows. Section 2 reviews physical monitoring, multi-parameter learning, CCAF images, physics-informed learning, and cross-domain evaluation. Section 3 specifies the dual-modality network and the channel-wise decision rule. Section 4 describes the dataset and baselines, then reports four evaluation protocols: within-scenario, within-dataset cross-scenario, cross-dataset hold-out, and ablation. Section 5 discusses the findings, limitations, and implications. Section 6 concludes.

2 Related Work

2.1 Single-Metric Physical Monitoring

SQM was introduced to measure distortion of the correlation peak at the early, late, and prompt taps, and was later used for spoofing detection. Classical Ratio, Delta, and ELP metrics fail quickly under frequency unlock or matched-power spoofing, which has led to two reinforcements. One is a temporal high-pass. Moving-variance SQM uses a short-window variance to emphasize peak jitter at takeover rather than the absolute level (Sun et al. 2018a), and a C/N_0 moving variance combined with an improved Ratio has been used for meaconing detection (Zhu et al. 2022). The other operates in metric space, either by a linear combination of several SQM statistics (Sun et al. 2018b) or by quadrature-branch SQM that reduces multipath false alarms (Sun et al. 2021). Frequency-domain processing can further reduce overlap with time-domain interference (Li S et al. 2025). Multicorrelator distortion monitoring still reduces the correlator profile to one distortion statistic (Shang et al. 2023).

These scalar tests still leave part of the attack uncovered, so later work reads the peak shape more directly. Power–distortion monitoring runs a power test and a distortion test together (Wesson et al. 2018). A spoofing advantage near 1dB leaves the power test nearly silent, so

the distortion test has to decide; a large advantage reverses that split. Small-delay multiple peaks and region-of-interest recognition on the correlation peak show, in the same direction, that the local code-delay neighborhood already contains separable geometry (Li et al. 2020; Yang et al. 2022). Ji et al. add two extra correlators, form a likelihood from the correlation residue, and pool the satellite tests into one receiver-level call; the TEXTBAT detection rate they report is high (Ji et al. 2024). The residue is used here in a different role: it enters a per-channel gate rather than a receiver-level likelihood. IKF-SQM takes a related step in the time domain. It smooths in-phase and quadrature SQM with an indirect Kalman filter and sets the threshold from the filter’s steady-state variance, which on TEXTBAT shortens the alert and reduces metric jitter (Jin et al. 2026b). The common limit is that each of these monitors still hinges on one family of peak statistics.

2.2 Multi-Parameter Learning and Unknown Scenarios

Because a single metric is limited, later detectors stack SQM, C/N_0 , pseudorange–Doppler consistency, clock bias and drift, and position residuals into multi-parameter vectors. A multi-parameter SVM is a milestone of this line and substantially raises F1 relative to a single SQM metric (Chen et al. 2022). Observation cross-correlation (Semanjski et al. 2019) and classification after Hilbert-envelope enhancement (Wang et al. 2024) belong to the same family. Tracking features that a conventional receiver can extract have been standardized, and both supervised classification and unsupervised representation learning have been reported (Iqbal et al. 2024a). An inductive framework uses a variational autoencoder or a generative adversarial network to describe the clean manifold and emphasizes separability of unseen attacks (Iqbal et al. 2024b). CGAN-ANN compresses the description to four core radio-frequency and tracking quantities and uses conditional generation to fill deformations of unknown scenarios (Chen L et al. 2025). A LightGBM detector in the acquisition domain can also reach a high detection rate on TEXTBAT and is often paired with array direction finding (Jin et al. 2025).

These methods show that the features are already informative. The remaining bottleneck is how the features are fused, and how the data are split. When the scenario appears in training, concatenation followed by an SVM can detect spoofing accurately. On an unknown scenario, the interval just after takeover and the matched-power segment are usually lost first, because the decision surface is fit to steady levels rather than to the few seconds in which the physical quantities begin to break. Generative methods define “unknown” as a departure in feature space. Related generalization tests instead train across receivers, front ends, or rewards: multi-receiver ensembles (Pervysheva et al. 2024), cross-dataset radio-frequency fingerprints (Marata et al. 2025), and deep reinforcement learning on tracking-loop states (Vemuri et al. 2026). They are effective for overpowered patterns that have never been seen. Jam-then-spoof interference remains outside that description: the channel is already false-locked, and the SQM metrics again resemble the clean distribution.

2.3 Correlator Planes and CAF Images

Once multi-correlator outputs are arranged on a code-delay–time plane, spoofing appears as peak splitting, translation, or asymmetric broadening. The bound recalled in Section 1 is why SpoofConvNet drops a fixed linear SQM combiner in favor of convolution on that grid, in a network sized for consumer devices (Jin et al. 2026a). Small-delay forwarding under multipath can also be recognized by convolution on correlator images (Jin et al. 2024). In the acquisition domain, CAF/CCAF images are more common: deep networks detect the number and location of peaks on CAF patches (Borhani-Darian et al. 2024), and a correlation-peak region of interest is used for interference recognition (Yang et al. 2022). A sparse-decomposition framework further separates authentic and counterfeit components in the code-tracking loop and, on TEXTBAT, supports detection, classification, and position recovery.

Two design choices follow from this line. Nonlinear processing of the correlator plane is preferred to a linear combination of Ratio values, so a convolutional encoder is retained and the heatmap is not collapsed to two or three scalars before classification. The local code-delay neighborhood is also sufficient to represent the secondary peak and the peak walk during takeover, so a full Doppler-delay search is not required. Local Mag/Q heatmaps are therefore used as the CCAF image. Both are peak-normalized, which keeps absolute power out of the image while preserving the relative height of a secondary peak.

2.4 Physics-Informed Learning

Physics-informed machine learning has so far been used for GNSS detection in three ways. A physical relation can be written into the loss: CCC-PCNNs place a code-phase versus carrier-phase first-difference penalty beside a one-dimensional convolutional classifier, and that penalty was reported to beat using the same quantity only as an extra input (Albright et al. 2020; Islam et al. 2024; Song et al. 2026). Residuals can instead be used as features for a boosted tree or a network. A physical model can also be used as a filter, as in IKF-SQM (Jin et al. 2026b).

CCC-PCNNs are the closest methodological neighbor, and their limitation is specific. Code-carrier consistency is a strong invariant of authentic GNSS signals, but it is not sensitive to every attack. Synchronized overpowered spoofing can rewrite power and automatic gain before the code-carrier relation is clearly broken. The jamming segment of combined interference contaminates carrier-phase continuity. In the middle of a matched-power interval, the first difference can fall back into the noise. Writing a single code-carrier penalty symmetrically on every sample also conflicts with the floor residual of a weak antenna or of multipath. More generally, a physics-informed or physics-constrained loss treats $r \approx 0$ on the clean state as a law. Most post-loop test statistics used for spoofing detection are not zero under the null. Code-carrier consistency is therefore retained only as one coordinate r_{ccc} of a residual bank, together with SQM, peak walk, a secondary peak, and a clock residual. The residual vector $\mathbf{r}$ is stop-gradient with respect to the network input and is used only to condition the two gates.

3 Materials and Methods

This section specifies PI-DMCANs. Loop test statistics from receiver post-processing are treated as stop-gradient anomaly coordinates. They do not enter the network as ordinary features. They condition the causal cross-attention that fuses the tracking sequence with the CCAF image. A one-sided score-residual consistency term then constrains the detection score, and the output is a per-channel anomaly score that can be evaluated under an NP rule. Fig. 1 shows the data flow, including the causal cross-attention block. Fig. 3 and Fig. 4 show the two encoders.




Fig. 1 Overall architecture of PI-DMCANs

3.1 Dual-Modality Observations

3.1.1 Tracking Features

Observations for the tracking features are taken from the tracking loop and the consistency module of the software receiver. Inside each detection window of length $T_w = 1\text{s}$, every satellite channel is organized as

$$\mathbf{S}_{n,k} \in \mathbb{R}^{C \times T_s} \quad (1)$$

The six channels are fixed as

$$\{\text{Ratio, Delta, ELP, C/N}_0, \text{std}(\text{C/N}_0), \text{CMC}\} \quad (2)$$

Ratio, Delta, and ELP are multi-correlator SQM metrics. They describe the early-late power ratio, early-minus-late deformation, and deformation of the early-plus-late power, respectively. C/N_0 and its within-window standard deviation describe the carrier-to-noise level and its short-term fluctuation. CMC is an absolute summary of the code-carrier consistency residual inside the window.

3.1.2 CCAF Image

For each channel, the local magnitude and quadrature of the cross-ambiguity function on a code-delay neighborhood of ± 10 chips are arranged as

$$\mathbf{H}_{n,k} \in \mathbb{R}^{2 \times T_h \times D}, \quad (3)$$

where T_h is the number of short frames, cropped causally at the window end, and D is the number of delay bins. Let $\mathbf{m}_t \in \mathbb{R}^D$ and $\mathbf{q}_t \in \mathbb{R}^D$ be the magnitude and quadrature profiles of frame t . Absolute brightness is removed, and the relative height of a secondary peak is retained, by peak-normalizing each frame:

$$\mathbf{m}_t = \frac{\mathbf{m}_t}{\|\mathbf{m}_t\|_\infty + \varepsilon}, \quad \mathbf{q}_t = \frac{\mathbf{q}_t}{\|\mathbf{q}_t\|_\infty + \varepsilon}, \quad (4)$$

where $\varepsilon > 0$ is a numerical stabilizer. The delay axis is measured in chips, so signal structures with different code rates fall on the same geometry. Both channels share the Mag peak as the scale, which preserves the height ratio of a secondary peak to the main peak. The two channels of $\mathbf{H}$ are $\{\mathbf{m}_t\}$ and $\{\mathbf{q}_t\}$.

Fig. 2 contrasts the two cases. On a clean channel the Mag profile is a sharp single peak, and the short-time heatmap forms a stable narrow ridge near zero delay. On a spoofed channel the same satellite shows two peaks separated by several chips, and the heatmap splits into two bright bands along the delay axis. This delay–time geometry is captured directly by the CCAF encoder.




Fig. 2 CCAF image. Left: a clean channel, with a single-peak Mag profile and a narrow zero-delay ridge. Right: a spoofed channel, with two peaks separated by several chips and two bright bands along the delay axis. Top: one peak-normalized Mag profile versus code delay. Bottom: the short-time Mag heatmap inside the window

3.2 Dual-Modality Encoders

The two encoders map the heterogeneous observations to a common embedding dimension $d_z = 48$ and are trained end to end with the fusion, the gate, and the detection head.

3.2.1 Tracking Encoder

$\mathbf{S}_{n,k} \in \mathbb{R}^{C \times T_s}$ is transposed into a length- T_s sequence of vectors $\{\mathbf{x}_t\}_{t=1}^{T_s}$, where $\mathbf{x}_t \in \mathbb{R}^C$ is the six-dimensional tracking vector in bin t . The tracking encoder E_{trk} is a one-layer bidirectional long short-term memory network (Hochreiter and Schmidhuber 1997; Graves and Schmidhuber 2005). In each direction, the forget gate, input gate, output gate, and cell state follow the standard update

$$\mathbf{f}_t = \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{b}_f), \quad (5)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{b}_i), \quad (6)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{b}_o), \quad (7)$$

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_{xc}\mathbf{x}_t + \mathbf{W}_{hc}\mathbf{h}_{t-1} + \mathbf{b}_c), \quad (8)$$

$$\mathbf{c}_t = \mathbf{f}_t \square \mathbf{c}_{t-1} + \mathbf{i}_t \square \tilde{\mathbf{c}}_t, \quad (9)$$

$$\mathbf{h}_t = \mathbf{o}_t \square \tanh(\mathbf{c}_t), \quad (10)$$

where σ is the element-wise sigmoid, $\square$ is the element-wise product, and $\mathbf{h}_t, \mathbf{c}_t \in \mathbb{R}^{d_{\text{hid}}}$ with $d_{\text{hid}} = 48$. Forward and backward hidden states are concatenated and passed through a two-layer linear map to obtain the tracking tokens

$$\mathbf{h}_t^{\text{bi}} = [\vec{\mathbf{h}}_t; \overleftarrow{\mathbf{h}}_t] \in \mathbb{R}^{2d_{\text{hid}}}, \quad \mathbf{Z}_{n,k}^{\text{trk}} = (F_{\text{trk}}(\mathbf{h}_t^{\text{bi}}))_{t=1}^{T_s} \in \mathbb{R}^{T_s \times d_z} \quad (11)$$

Inside a closed 1s window, the bidirectional recurrence uses both past and future bins, so that the short nonstationarity around takeover can be represented. This complements the causal recurrence used on the CCAF branch. The data path is shown in Fig. 3.




Fig. 3 Tracking encoder E_{trk}

3.2.2 CCAF Encoder

The two channels of $\mathbf{H}_{n,k} \in \mathbb{R}^{2 \times T_h \times D}$ are Mag and Q, with time length T_h and delay length D . The CCAF encoder E_{img} first extracts frame-level geometry with two two-dimensional convolutions on the local delay–time neighborhood (LeCun et al. 1998):

$$\mathbf{Y}^{(\ell)} = \text{ReLU}(\mathbf{W}^{(\ell)} * \mathbf{Y}^{(\ell-1)} + \mathbf{b}^{(\ell)}), \quad \ell = 1, 2, \quad (12)$$

where $\mathbf{Y}^{(0)} = \mathbf{H}_{n,k}$ and $*$ denotes two-dimensional convolution. In the implementation the channel counts are $2 \rightarrow 16 \rightarrow 32$, the kernel is 3×5 , and same padding is used. Adaptive max-pooling then reduces the delay axis to width 8. Channels and delay are flattened to a frame sequence $\mathbf{u}_t \in \mathbb{R}^{256}$ ($t = 1, \dots, T_h$).

The frame sequence is aggregated in time by a one-layer unidirectional gated recurrent unit (Cho et al. 2014), so that the temporal evolution of the peak shape is represented. With reset gate $\mathbf{r}_t$, update gate $\mathbf{z}_t$, and candidate state $\mathbf{h}_t$, the recurrence is

$$\mathbf{z}_t = \sigma(\mathbf{W}_z\mathbf{u}_t + \mathbf{U}_z\mathbf{h}_{t-1} + \mathbf{b}_z), \quad (13)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r\mathbf{u}_t + \mathbf{U}_r\mathbf{h}_{t-1} + \mathbf{b}_r), \quad (14)$$

$$\mathbf{h}_t = \tanh(\mathbf{W}_h\mathbf{u}_t + \mathbf{U}_h(\mathbf{r}_t \square \mathbf{h}_{t-1}) + \mathbf{b}_h), \quad (15)$$

$$\mathbf{h}_t = (1 - \mathbf{z}_t) \square \mathbf{h}_{t-1} + \mathbf{z}_t \square \mathbf{h}_t. \quad (16)$$

The hidden states are the CCAF tokens

$$\mathbf{Z}_{n,k}^{\text{img}} = (\mathbf{h}_t)_{t=1}^{T_h} \in \mathbb{R}^{T_h \times d_z} \quad (17)$$

The data path is shown in Fig. 4. After two 3×5 convolutions and delay-axis pooling, a unidirectional GRU aggregates the Mag/Q heatmap along causal time.




Fig. 4 CCAF encoder E_{img}

3.3 Dual-Modality Fusion

The two modalities do not share a time scale or a missingness pattern, so they are not concatenated. CCAF tokens are first retrieved by the tracking tokens, which aligns a CCAF context to the tracking time axis. Residual-conditioned gates then set the two stream weights, and the fused representation is mapped to a channel score.

3.3.1 Causal Cross-Attention

Tracking tokens are the queries, and CCAF tokens are the keys and values. For multi-head scaled dot-product attention, head h is

$$\mathbf{Q}_h = \mathbf{Z}^{\text{trk}} \mathbf{W}_h^Q, \quad \mathbf{K}_h = \mathbf{Z}^{\text{img}} \mathbf{W}_h^K, \quad \mathbf{V}_h = \mathbf{Z}^{\text{img}} \mathbf{W}_h^{V'}, \quad (18)$$

$$\text{Attn}_h = \text{softmax} \left(\text{mask} \left(\frac{\mathbf{Q}_h \mathbf{K}_h^T}{\sqrt{d_h}} \right) \right) \mathbf{V}_h. \quad (19)$$

The heads are concatenated and linearly projected to the CCAF context

$$\mathbf{C}_{n,k} = \text{Concat}_h(\text{Attn}_h) \mathbf{W}^O \in \mathbb{R}^{T_s \times d_z} \quad (20)$$

Useful evidence is not always in the same modality. Overpowered spoofing often rewrites power first, matched-power spoofing depends more on the peak shape, and after false lock the Ratio can return close to the clean value. If the two streams are added with equal weight, a clean window is easily biased by CCAF noise, and a difficult case is diluted by the stream that has not changed. The two weights are therefore set by a physical residual bank.

3.3.2 Physical Residual Bank

Five representative physical test statistics are extracted during receiver processing. Their excursions relative to a clean reference form the residual bank. For each detection window a five-dimensional raw statistic is constructed,

$$\tilde{\mathbf{r}} = [\tilde{r}_{\text{ccc}}, \tilde{r}_{\text{sqm}}, \tilde{r}_{\text{warp}}, \tilde{r}_{\text{twin}}, \tilde{r}_{\text{clk}}]^T \quad (21)$$

$\tilde{r}_{\text{ccc}}$ measures code-carrier consistency; $\tilde{r}_{\text{sqm}}$ measures the departure of the early-late power ratio and its short-term jitter from the clean level; $\tilde{r}_{\text{warp}}$ measures the code-chip offset of the local CAF main peak from zero delay; $\tilde{r}_{\text{twin}}$ measures a far-sidelobe secondary peak; and $\tilde{r}_{\text{clk}}$ measures anomalous receiver clock drift. Heavy-tailed quantities are compressed by $\psi(x) = \log(1 + |x|)$, so that a code-minus-carrier residual of hundreds of meters does not dominate the other coordinates. When peak-walk and secondary-peak residuals are later centered, the absolute deviation from the clean median is used. Both overpowered main-peak collapse and matched-power peak splitting or translation are two-sided excursions, so those

two statistics are two-sided. Code-carrier, SQM, and clock residuals remain one-sided upper excursions.

The code-carrier residual is

$$\tilde{r}_{\text{ccc}} = \max(\psi(|\varepsilon_{\text{CMC}}|), \psi(|\varepsilon_{\text{CCD}}|)) \quad (22)$$

Here ε_{CMC} is the code-minus-carrier residual and ε_{CCD} is the code-carrier Doppler consistency residual. The larger of the two compressed values is kept, so that either a significant code-phase divergence or a Doppler inconsistency is sufficient to trigger the coordinate. This coordinate follows the enhanced code and carrier Doppler consistency test used to detect time-synchronization attacks (Feng et al. 2024). Let f_c and f_{code} be the carrier and code reference frequencies, $f_{\text{D,code}}$ the code Doppler, $f_{\text{D,carr}}$ the carrier Doppler, and c the speed of light. The code-carrier Doppler residual, converted to range, is

$$\varepsilon_{\text{CCD}} = \frac{c}{f_{\text{code}}} \cdot \left(\frac{f_{\text{D,carr}}}{f_c / f_{\text{code}}} - f_{\text{D,code}} \right) \quad (23)$$

Because f_c and f_{code} are those of the constellation, channels with different code rates share the same $\tilde{r}_{\text{ccc}}$.

The SQM residual is

$$\tilde{r}_{\text{sqm}} = \left| \text{Ratio} - \mu_{R,c} \right| + \frac{1}{2} \text{std}(\text{Ratio}) \quad (24)$$

where $\mu_{R,c}$ is the median Ratio of clean training windows in constellation family c , and $\text{std}(\text{Ratio})$ is the within-window standard deviation of Ratio. The two terms capture a level offset and short-term jitter. Because the signal structure differs across constellations, a shared median would cause misclassification, so $\mu_{R,c}$ is estimated separately for each family.

The peak-walk residual is

$$\tilde{r}_{\text{warp}} = |d^{\hat{a}}|, \quad d^{\hat{a}} = \arg \max_d |m(d)|, \quad (25)$$

where $m(d)$ is the local CAF magnitude profile and d is the delay in chips relative to zero. The code-chip offset of the main peak is computed on each of the T_h frames, and the median is retained. During matched-power pull-off the main peak leaves zero delay, so $|d^{\hat{a}}|$ measures that displacement. Under a clean lock the main peak remains near zero delay and the term is close to zero.

The secondary-peak residual is

$$\tilde{r}_{\text{twin}} = T \quad (26)$$

Let the peak-normalized profile be $\bar{\mathbf{m}} = \mathbf{m} / \|\mathbf{m}\|_{\infty}$. Only when zero delay itself is a local peak and $\bar{m}(0) \geq 0.80$ is T set to the normalized height of the second local maximum at $|d| \geq 2$ chips; otherwise $T = 0$ for that frame.

The clock-drift residual is

$$\tilde{r}_{\text{clk}} = |\dot{b}_u|, \quad (27)$$

where $\dot{b}_u$ is the receiver clock drift.

Let $m_{c,j}$ and $\rho_{c,j}$ be the median and the empirical 90th percentile of component j on clean training windows of constellation family c . For the peak-walk and secondary-peak

components the absolute deviation from the median is taken; the other components use $\tilde{r}_j$ directly. A one-sided excursion above the clean 90th percentile is then formed:

$$\delta_j = \begin{cases} |\tilde{r}_j - m_{c,j}|, j \in \{\text{warp}, \text{twin}\} \\ \tilde{r}_j, \text{otherwise} \end{cases} \quad (28)$$

$$r_j = \text{clip}(\text{ReLU}(\delta_j - \rho_{c,j}), 0, r_{\max})$$

The percentile vector $\mathbf{p}_c$ is estimated by constellation family. The 90th percentile leaves margin for clean fluctuations such as multipath and weak satellites. The ReLU keeps only the positive excess, and the clip prevents cycle slips and momentary loss of lock from dominating the gate.

3.3.3 Gating Network

A shallow multilayer perceptron maps $\mathbf{r}$ to a two-dimensional gate logit, and a learnable soft physical prior is added. From the information geometry of the residuals, the tracking gate is biased toward $\tilde{r}_{\text{ccc}}$, $\tilde{r}_{\text{sqm}}$, and $\tilde{r}_{\text{clk}}$, and the CCAF gate toward $\tilde{r}_{\text{warp}}$, $\tilde{r}_{\text{twin}}$, and $\tilde{r}_{\text{sqm}}$:

$$\begin{aligned} \ell_{\text{trk}}^{\text{pri}} &= a_{\text{ccc}} r_{\text{ccc}} + a_{\text{sqm}} r_{\text{sqm}} + a_{\text{clk}} r_{\text{clk}} \\ \ell_{\text{img}}^{\text{pri}} &= b_{\text{warp}} r_{\text{warp}} + b_{\text{twin}} r_{\text{twin}} + b_{\text{sqm}} r_{\text{sqm}} \end{aligned} \quad (29)$$

$$\mathbf{g}_{n,k} = [\mathbf{g}_{\text{img}}, \mathbf{g}_{\text{trk}}] = \sigma(\text{MLP}(\mathbf{r}_{n,k}) + \kappa[\ell_{\text{img}}^{\text{pri}}, \ell_{\text{trk}}^{\text{pri}}]) \in (0, 1) \quad (30)$$

The coefficients a_i and b_i are nonnegative and equal to a softplus of a free parameter. They are trained end to end with the perceptron. The prior scale is $\kappa = 0.35$, so that the prior remains weaker than the free perceptron and the modality assignment is not a hard rule. The gate values fuse the two streams, followed by a pre-normalization residual block (Vaswani et al. 2017):

$$\mathbf{U}_{n,k} = \mathbf{g}_{\text{trk}} \mathbf{Z}_{n,k}^{\text{trk}} + \mathbf{g}_{\text{img}} \mathbf{C}_{n,k}, \quad (31)$$

$$\mathbf{F}_{n,k} = \text{LN}(\mathbf{U}_{n,k}) + \text{DropPath}(\text{FFN}(\text{LN}(\mathbf{U}_{n,k}))). \quad (32)$$

Layer normalization standardizes $\mathbf{U}$ along the feature axis and limits scale drift between the two streams. The position-wise feed-forward network mixes channels at each time step and supplies pointwise interactions that cross-attention does not cover. Stochastic depth is applied only on the residual branch (Huang et al. 2016), which reduces overfitting to a single modality.

Attention pooling over time yields the channel embedding

$$\tilde{\mathbf{z}}_{n,k} = \sum_{t=1}^{T_s} \alpha_t \mathbf{F}_{n,k}[t], \quad \mathbf{a} = \text{softmax}(\mathbf{F}_{n,k} \mathbf{q}), \quad (33)$$

where $\mathbf{q}$ is a learnable query. A shallow detection head maps the embedding to a channel-level anomaly score

$$s_{n,k} = f_{\text{det}}(\tilde{\mathbf{z}}_{n,k}) \in \mathbb{R}, \quad (34)$$

where $f_{\text{det}}: \mathbb{R}^{d_z} \rightarrow \mathbb{R}$ is a linear layer.

3.4 Training Objective

Whether a transferable physical representation is learned depends on where the training objective places its gradients. The total objective is the sum of a detection loss, a gate regularizer, a channel contrastive loss, and a score-residual consistency term:

$$\mathcal{L} = \mathcal{L}_{\text{det}} + \lambda_g \mathcal{L}_{\text{gate}} + \lambda_{\text{nce}} \mathcal{L}_{\text{nce}} + \lambda_{\text{cons}} \mathcal{L}_{\text{cons}} \quad (35)$$

The detection loss supervises $s_{n,k}$ directly, because the primary task is channel-level binary classification. Standard cross-entropy is dominated by the numerous clean windows that are already confidently classified, which dilutes the gradient on difficult cases. A class-weighted focal loss is therefore used. With channel label $y_{n,k} \in \{0,1\}$, the predicted probability is

$$\hat{p}_{n,k} = \sigma(s_{n,k}) = \frac{1}{1 + e^{-s_{n,k}}}, \quad (36)$$

and the probability of the correct class is

$$p_t = \begin{cases} \hat{p}_{n,k}, & y_{n,k} = 1, \\ 1 - \hat{p}_{n,k}, & y_{n,k} = 0. \end{cases} \quad (37)$$

Following the focal loss of Lin et al. (2017), easy and hard samples are reweighted on top of a positive-class-weighted binary cross-entropy:

$$L_{\text{det}} = \mathbb{E}[w_{n,k} \alpha_y (1 - p_t)^{\gamma_n} (-\log p_t)], \quad (38)$$

where $\gamma_n = 2$ and α_y balances the positive and negative classes. The sample weight $w_{n,k}$ raises the positive class around a spoofing takeover with a Gaussian kernel of width about $5s$, because the few seconds before and after takeover are where the physical trace first appears and where channel-level misses concentrate.

The gate is the main path by which residuals enter the fusion. If only L_{det} is used, g can collapse to a constant independent of $\mathbf{r}$, and residual conditioning is lost. Binary entropy is largest at $g = 0.5$, so an entropy penalty pushes the gate away from the neutral region. Larger residuals should force a firmer choice of modality, which is implemented by the weight $\omega(\mathbf{r})$. A sum floor is also imposed so that both gates are not driven too small and the fusion does not collapse. For a scalar gate $g \in (0,1)$, the binary entropy is

$$H_b(g) = -g \log g - (1 - g) \log(1 - g) \quad (39)$$

The gate regularizer is

$$L_{\text{gate}} = \mathbb{E}[\omega(\mathbf{r}) \frac{1}{2} (H_b(g_{\text{img}}) + H_b(g_{\text{trk}}))] + \eta \mathbb{E}[\max(0, \xi - g_{\text{img}} - g_{\text{trk}})] \quad (40)$$

where $\omega(\mathbf{r}) = \|\mathbf{r}\|_1 / (\|\mathbf{r}\|_1 + \varepsilon)$ is only a sample weight, $\eta = 0.05$, and $\xi = 0.35$.

The consistency term asks the score to point in the same direction as an already observed above-percentile residual. It does not require the network to drive $\mathbf{r}$ to zero, which is the role of the symmetric code-carrier penalty in CCC-PCNNs (Song et al. 2026). A symmetric penalty would punish the multipath or weak-satellite floor of clean states, and it fits poorly in the middle of a matched-power interval, where the residual is small. The term used here is therefore one-sided. When the residual is already large and the model still calls the window clean, $(1 - \hat{p})\|\mathbf{r}\|_1$ increases and the gradient raises s . Clock drift is handled separately because it is a receiver common mode: the same clock error raises every constellation channel at once, and a channel-wise score is only weakly sensitive to it. Stronger detection of common-mode time pushes such as TEXBAT ds7 typically requires an explicit clock-drift or timing integrity monitor, and possibly an inertial unit or an external time reference. The clock coordinate $\tilde{r}_{\text{clk}}$ is therefore excluded from the consistency loss and is used only for tracking stream gating. With $\mathbf{r}^{\hat{a}} = [r_{\text{ccc}}, r_{\text{sqm}}, r_{\text{warp}}, r_{\text{twin}}]^T$,

$$L_{\text{cons}} = \mathbb{E}[(1 - \hat{p}_{n,k}) \|\mathbf{r}^{\hat{a}}\|_1] \quad (41)$$

The detection target is channel-level. If anomalous windows of different satellites are not distinguished, they can be grouped together because they share a front end, and the model then learns a scenario or a device rather than the physical state of that satellite channel. A contrastive loss is therefore applied to the channel embedding. Let $\mathbf{z}_i$ be the ℓ_2 -normalized embedding of $\tilde{\mathbf{z}}$. A multi-positive InfoNCE loss is used. The positive set $P(i)$ contains other windows of the same scenario and the same PRN, so that the embedding encodes the physical state of that channel rather than the window center time. The negative set $N(i)$ contains anomalous windows of different PRNs, so that several channels are not collapsed to one embedding merely because they share a front end under combined interference:

$$L_{\text{nce}} = -E_i \log \frac{\sum_{j \in P(i)} \exp(\mathbf{z}_i \mathbf{z}_j / \tau_{\text{nce}})}{\sum_{j \in P(i) \cup N(i)} \exp(\mathbf{z}_i \mathbf{z}_j / \tau_{\text{nce}})}, \quad (42)$$

where τ_{nce} is a temperature.

3.5 NP Threshold and Channel Output

The channel score in Equation (34) is first calibrated on the training sample by an affine map $s \leftarrow as + b$. On cross-scenario leave-one-scenario-out, the main tables and the five baselines share an empirical percentile of the clean training scores, so that the decision on a held-out scenario does not use test labels:

$$\tau = \text{Quantile}_{1-\alpha}(\{s_{n,k} : (n,k) \in H_0^{\text{tr}}\}), \quad \alpha = 0.05, \quad (45)$$

where $\alpha \in (0,1)$ is the target false-alarm level and H_0^{tr} is the set of clean training channels. $\text{Quantile}_q(S)$ denotes the empirical q -quantile of S . The same train-clean percentile is used under both the within-scenario and the cross-scenario protocols. After the threshold is moved to a held-out scenario the false-alarm rate may depart from the nominal 5%, but the operating point remains comparable across methods.

4 Results

Channel-level spoofing detection was evaluated under four complementary protocols: (i) within-scenario evaluation, which tests discrimination when the attack scenario is known; (ii) within-dataset cross-scenario leave-one-scenario-out, which holds out one scenario while training only on other scenarios from the same dataset; (iii) cross-dataset leave-one-scenario-out on a joint TEXTBAT–MuCoGS pool, which tests generalization when the held-out scenario is unseen and training may draw on the other dataset; and (iv) ablation of the residual and fusion components. All protocols use the train-clean 95th-percentile Neyman–Pearson threshold, and the same comparison against five baselines.

4.1 Datasets

Channel-level spoofing detection was evaluated on three raw radio-frequency corpora that are public or have been publicly described: TEXTBAT (Humphreys et al. 2012), OAKBAT (Albright et al. 2020), and the multi-constellation spoofing set MuCoGS (Chen et al. 2026). The three corpora differ in front-end sampling rate, constellation mix, and spoofing type. A compact comparison is given in Table 1. All features were extracted from baseband I/Q by the open-source software receiver family FGI-GSRx (Islam et al. 2024), so that the methods are compared on the same tracking-domain observables. A chi-square test on raw samples from the same receiver reports interference without a per-satellite decision (Bhuiyan et al. 2026).

Item	TEXBAT	OAKBAT	MuCoGS
Recording year	2012	2020	2026
Constellation	GPS L1	GPS L1, Galileo E1	GPS L1, Galileo E1, BDS B1I
Scenarios	ds2–ds7	os2–os6	ms1–ms6
Attack types	synchronized overpowered / matched-power	synchronized overpowered / matched-power	position / time / meaconing and combined interference
Sampling rate	25MHz	5MHz	62.5MHz
Receiver motion	static / dynamic	static / dynamic	static

Table 1 Datasets used for evaluation

TEXBAT. The Texas Spoofing Test Battery is an early standard dataset for baseband GPS L1 spoofing evaluation. It contains clean controls together with injection scenarios that use several power advantages and pull-off strategies (Humphreys et al. 2012). Scenario ds1 is not a realistic operating case and was excluded, so only ds2–ds7 enter the evaluation. The clean side and the spoofed side each last 1200s. Clean labels come from two places. The first 300s of cleanStatic and of cleanDynamic were taken, and every satellite channel in those segments was labeled clean ($y = 0$). The 100s before injection in ds2–ds7 were labeled clean as well. The spoofed side is the 200s after injection in the same files ($y = 1$). The clean control matched to the motion state is then added when a fold is held out. For a static fold (ds2, ds3, ds4, or ds7) the added windows come from cleanStatic; for a dynamic fold (ds5 or ds6) they come from cleanDynamic. The test set therefore includes a clean control, and is not limited to the prefix of the attack file.

OAKBAT. OAKBAT is a simulator-generated companion to TEXBAT. Its GPS L1 C/A files follow the key settings of TEXBAT ds1–ds6, and a parallel Galileo E1 collection repeats that scenario design (Albright et al. 2020). Only os2–os6 are used here. Those files have no separate clean recording, so the label is taken from the scenario of each attack file: the 100s before injection are clean, and the 100s after injection are spoofed.

MuCoGS. Public corpora contain few combined-interference scenarios, so a multi-constellation collection platform was built for this study. The resulting dataset, MuCoGS, provides raw I/Q of clean and attacked channels on a scenario synchronized with live sky signals (Chen et al. 2026). Live sky signals were received by a survey antenna and recorded by a software radio at a center frequency of 1575.42MHz, with an acquisition bandwidth of about 50MHz and a sampling rate of 62.5MHz. That bandwidth covers GPS L1, Galileo E1, and BeiDou B1I/B1C in one recording. Seven scenario classes are publicly described and are denoted ms1–ms7 (Table 2): direct position spoofing, time spoofing, and meaconing, together with jam-then-spoof combinations of the same three types. The jamming segment is a swept interference, with a sweep period of about 1s and a bandwidth of about 20MHz. Spoofing injection starts about 10s before the jammer is switched off, so that after reacquisition the receiver is already false-locked. The dataset is hosted by the Science Data Bank. Each record has two paths. CH0 is the attacked antenna, and CH1 is a control antenna synchronized with the sky signal.

Labels are assigned jointly by antenna path and by scenario, and the two collection channels are taken over the same time window. Scenarios ms1–ms6 were used, and the extracted features cover GPS L1, BeiDou B1I, and Galileo E1. The window that enters the detector depends on the intrusion type. Scenarios ms1 and ms3 are direct intrusion with a slow pull-

off, and the 100s after spoofing injection were extracted. Scenarios ms2, ms4, and ms6 are combined intrusion. The jamming interval is omitted, and only the 100s after spoofing injection are kept, which matches a receiver that has lost lock and then reacquired. Scenario ms5 is direct meaoning, and the 60s after injection were extracted. Samples from CH0 were labeled spoofed and samples from CH1 were labeled clean, so the two classes are balanced.

ID	Intrusion	Spoof type	Trajectory	Duration (s)
ms1	direct	position	pull-off	180
ms2	combined (jam → spoof)	position	jump	180
ms3	direct	time	pull-off	180
ms4	combined (jam → spoof)	time	jump	180
ms5	direct	meaoning	jump	120
ms6	combined (jam → spoof)	meaoning	jump	180
ms7	—	spoofers only	static	120

Table 2 Scenario summary for MuCoGS

4.2 Baseline Detectors

Five baselines were reimplemented under the same window centers, channel labels, and NP threshold, so that the gain of dual-modality fusion and residual gating can be isolated. Unless noted otherwise, sequence models take the tracking features in Equation (2) as input.

1. **BiLSTM.** A bidirectional long short-term memory network on the channel sequence of Equation (2), used to measure how well a purely temporal memory model holds up under a within-scenario split and under cross-scenario hold-out (Mao et al. 2025).
2. **Transformer.** A light self-attention encoder on the same short tracking sequence, used as a control for global temporal dependence (Hao et al. 2026).
3. **Multi-parameter SVM.** The within-window sequence is compressed to statistical vectors such as the mean and the standard deviation, and an SVM is trained. This baseline represents concatenated multi-parameter classification (Iqbal et al. 2024a).
4. **CAF-CNN.** A convolutional network whose only input is the local Mag/Q code-delay–time heatmap. It represents CAF image detection and measures performance when only the CCAF image is available (Jin et al. 2024).
5. **CCC-PCNNs.** A code–carrier consistency physics-constrained network: a physical term based on code–carrier consistency is applied in parallel with a one-dimensional convolutional classification loss, and the composite decision of the source paper is used (Song et al. 2026). The feature path was reimplemented from the published setting and is treated as the most recent physics-informed baseline.

4.3 Within-Scenario Evaluation

Each scenario was trained and evaluated on its own, so that the methods are compared when the attack scenario is shared by the training and test windows. A label-stratified random split was used: 70% for training, of which 10% of the training pool was held out for validation, and 30% for testing. Because the test windows belong to the same scenario, the decision surface can use the steady level of that spoofing pattern. The split confirms that the two modalities are separable, and it supplies the upper reference for the cross-scenario retention rate. In every fold the threshold is the 95th percentile of the clean training scores of each constellation, the same rule used for the cross-scenario protocol and for the five baselines. The two protocols share the network and loss settings in Table 3 and differ only in the split.


Group	Parameter	Value
Network	embedding dimension d_z	48
	cross-attention heads	4
	gate-prior scale κ	0.35
	λ_{cons}	0.50
Loss weights	λ_{ncc}	0.03
	λ_g	0.12
	focal-loss exponent γ_{fl}	1

Table 3 Experimental parameter settings

Table 4 reports the mean Accuracy over 17 folds. Fig. 5 shows the corresponding mean receiver operating characteristic (ROC). Panel (a) averages all 17 folds. Panels (b)–(d) give the within-dataset means for TEXTBAT, OAKBAT, and MuCoGS. For each fold an ROC was computed from the test scores, and the detection probability P_D was then averaged on a common false-alarm grid. When the scenario is known, PI-DMCANs leads the within-scenario table: Accuracy is 0.965 over the 17 folds, ahead of CCC-PCNNs at 0.941 and CAF-CNN at 0.940. The SVM follows at 0.924, then the Transformer at 0.865 and the BiLSTM at 0.840. The six MuCoGS folds give 0.966 for PI-DMCANs and 0.933 for CAF-CNN. Split by constellation, the PI-DMCANs means are 0.989 on BDS B1I, 0.980 on GAL E1, and 0.920 on GPS L1.

Method	Accuracy	TEXTBAT	OAKBAT	MuCoGS
SVM	0.924	0.911	0.957	0.909
CAF-CNN	0.940	0.917	0.974	0.933
Transformer	0.865	0.837	0.947	0.823
CCC-PCNNs	0.941	0.931	0.964	0.931
BiLSTM	0.840	0.834	0.936	0.766
PI-DMCANs	0.965	0.959	0.972	0.966

Table 4 Within-scenario channel-level Accuracy






Fig. 5 Mean ROC under the within-scenario protocol, computed from channel-level scores. (a) Mean of 17 folds. (b) TEXTBAT. (c) OAKBAT. (d) MuCoGS

4.4 Within-Dataset Cross-Scenario Evaluation

Within each dataset, one scenario was held out for testing and the remaining scenarios of that dataset only were used for training. No fine-tuning is performed at test time. This protocol isolates within-dataset generalization to an unseen scenario without borrowing examples from another front end.

Table 5 reports the mean Accuracy under the same train-clean NP threshold. PI-DMCANs reaches a mean Accuracy of 0.928 (TEXTBAT 0.928, OAKBAT 0.955, MuCoGS 0.906), ahead of CAF-CNN at 0.893 and CCC-PCNNs at 0.866. Relative to the within-scenario mean of 0.965 on the same folds, the within-dataset retention is 0.962.

Method	Accuracy	TEXTBAT	OAKBAT	MuCoGS
SVM	0.817	0.763	0.965	0.747
CAF-CNN	0.893	0.849	0.955	0.886
Transformer	0.863	0.801	0.963	0.841
CCC-PCNNs	0.866	0.797	0.945	0.868
BiLSTM	0.843	0.773	0.950	0.824
PI-DMCANs	0.928	0.928	0.955	0.906

Table 5 Within-dataset leave-one-scenario-out Accuracy

Fig. 6 expands the same NP operating point into a heatmap over the 17 scenarios. On the six TEXTBAT folds and the six MuCoGS folds, PI-DMCANs is the highest or tied for the highest. On OAKBAT the sequence baselines are already close to saturation, except that the os6 scenario is only 0.874.




Fig. 6 Heatmap of accuracy for each scenario under within-dataset cross-scenario evaluation

4.5 Cross-Dataset Leave-One-Scenario-Out Evaluation

In each fold one entire scenario was held out for testing and the remaining scenarios were used for training, so that the detector is scored on an attack it has not seen. The attack waveform and the radio-frequency scenario of the test fold do not appear in training, and no fine-tuning is performed at test time.

Table 6 compares the SVM, CAF-CNN, the Transformer, CCC-PCNNs, the BiLSTM, and PI-DMCANs on Accuracy, Precision, Recall, F1, and AUC. As we can see, the Accuracy, F1, and AUC of PI-DMCANs are higher than those of the five baselines. The margin over the sequence models is largest on TEXTBAT and on MuCoGS, where those models fall further behind the image-based detectors.

Method	Accuracy	Precision	Recall	F1	AUC
SVM	0.718	0.815	0.414	0.536	0.814
CAF-CNN	0.881	0.831	0.894	0.855	0.946
Transformer	0.850	0.841	0.753	0.783	0.910
CCC-PCNNs	0.805	0.821	0.724	0.736	0.887
BiLSTM	0.840	0.841	0.734	0.773	0.898
PI-DMCANs	0.932	0.904	0.922	0.908	0.979

Table 6 Channel-level comparison under cross-dataset evaluation

Fig. 7 also expands the same NP operating point into a heatmap over the 17 scenarios. On the six TEXTBAT folds and the six MuCoGS folds, PI-DMCANs is the highest or tied for the highest. On OAKBAT the sequence baselines are already close to saturation, and PI-DMCANs remains in the range 0.965–0.993.




Fig. 7 Heatmap of accuracy for each scenario under cross-dataset evaluation

Generalization can be summarized by the accuracy retention rate

$$ARR = \frac{\text{Accuracy}_{\text{cross}}}{\text{Accuracy}_{\text{within}}} \quad (46)$$

where the numerator and the denominator are means over the same set of folds. Leave-one-scenario-out is the primary measure of channel-level detection on unseen attack families. The within-scenario Accuracy is the denominator of ARR and the calibrated upper reference.

Table 7 converts those two protocols into an accuracy retention rate, by dataset and, inside MuCoGS, by constellation. Under the same threshold, the cross-scenario 17-fold mean Accuracy is 0.932, the within-scenario mean is 0.972, and ARR is 0.959. Most of the drop sits in TEXBAT, where ARR is 0.913, and in individual MuCoGS constellations. OAKBAT retains the within-scenario level, with ARR equal to 1.008.

Dataset	Accuracy (within)	Accuracy (cross)	ARR
TEXBAT	0.967	0.883	0.913
OAKBAT	0.973	0.981	1.008
MuCoGS	0.976	0.942	0.965
Mean	0.972	0.932	0.959

Table 7 ARR of PI-DMCANs, cross-dataset relative to within-scenario

4.6 Ablation Study

To isolate the contribution of each modality and of residual gating, four fusion variants were trained. Full (A0) is the residual-gated dual-modality monitor. Trk-only (A1) drops the CCAF encoder. CCAF-only (A2) drops the tracking encoder. Equal-fuse (A3) keeps both streams but fixes $g_{\text{trk}} = g_{\text{img}} = 0.5$, so the residual bank does not condition the fusion.

Tables 8–10 report Accuracy, P_D , P_{FA} and AUC on the within-scenario, within-dataset, and cross-dataset protocols. Under the within-scenario split, Full, CCAF-only, and Equal-fuse all saturate near 0.965. Trk-only falls to 0.912, mostly on MuCoGS, so the Mag/Q heatmap already carries most of the matched-scenario separability and the tracking stream alone is not enough on multi-constellation jam-then-spoof folds.

On within-dataset evaluation, Full remains at 0.928. Removing residual gating (Equal-fuse) changes Accuracy to 0.922. The single-modality drops are larger: Trk-only 0.893 and CCAF-only 0.840. The TEXTBAT subset concentrates most of the CCAF-only loss, consistent with matched-power cases in which the heatmap returns close to the clean ridge and the tracking stream is needed.

On the cross-dataset evaluation, Full scores 0.941. Equal-fuse scores 0.909, while Trk-only and CCAF-only fall to 0.860 and 0.821. Thus dual-modality fusion is required for cross-dataset hold-out, and residual-conditioned gating adds a further margin once both streams are present.

ID	Variant	Acc	P_D	P_{FA}	AUC
A0	Full	0.966	0.997	0.058	0.999
A1	Trk-only	0.912	0.888	0.060	0.964
A2	CCAF-only	0.965	0.996	0.058	0.999
A3	Equal-fuse	0.965	0.997	0.058	0.999

Table 8 Ablation under the within-scenario protocol

ID	Variant	Acc	P_D	P_{FA}	AUC
A0	Full	0.928	0.972	0.111	0.989
A1	Trk-only	0.893	0.824	0.066	0.949
A2	CCAF-only	0.840	0.939	0.229	0.947
A3	Equal-fuse	0.922	0.964	0.116	0.983

Table 9 Ablation under within-dataset protocol

ID	Variant	Acc	P_D	P_{FA}	AUC
A0	Full	0.941	0.969	0.075	0.988
A1	Trk-only	0.860	0.766	0.076	0.902
A2	CCAF-only	0.821	0.905	0.242	0.932
A3	Equal-fuse	0.909	0.920	0.100	0.970

Table 10 Ablation under cross-dataset protocol

5 Discussion

5.1 Within-Scenario Separability

Under the within-scenario protocol, the tracking features and the CCAF image still separate clean and spoofed channels, but the ranking is no longer tied between PI-DMCANs and CAF-CNN. The mean ROC in Fig. 5 remains clearly stratified. Table 4 gives Accuracy 0.965 for PI-DMCANs, 0.941 for CCC-PCNNs, 0.940 for CAF-CNN, and 0.924 for the SVM. On MuCoGS, PI-DMCANs reaches 0.966 while the SVM reaches 0.909. This within-scenario result is the calibrated upper reference for the cross-scenario and within-dataset hold-outs.

5.2 Hold-Out Generalization Relative to Baselines

Once the attack scenario is held out, the 17-fold Accuracy of CAF-CNN falls to 0.881, while PI-DMCANs remains at 0.932. The gap is consistent with the working hypothesis that residual-conditioned dual-modality fusion transfers more reliably than a single modality on an unseen scenario. Two mechanisms account for most of it.

On the difficult TEXTBAT folds, the tracking stream still carries a cue when the peak shape is close to the clean distribution. For matched-power ds4, PI-DMCANs remains at 0.933 and

CAF-CNN at 0.883, because code-carrier and SQM residuals continue to move after the heatmap has settled. Scenario ds7 is a harder case of the same family: a matched-power, carrier-phase-aligned time push. After aligned takeover, the power advantage and the SQM and peak-shape distortion stay close to the clean distribution for a long interval. Under leave-one-scenario-out every method is substantially lower, the scores lie close together, and missed detections, rather than false alarms, are the bottleneck. That pattern indicates limited separability of this attack family on an unseen scenario. Stronger detection of such common-mode time pushes typically requires an explicit clock-drift or timing integrity monitor (Feng et al. 2024; Li J et al. 2025), and possibly an inertial unit or an external time reference. Clock error is receiver-common, so a channel-wise score is only weakly sensitive to it. For that reason $\tilde{r}_{\text{clk}}$ is kept out of $\mathbf{L}_{\text{cons}}$ and is used only for tracking-stream gating.

The same fusion also explains the MuCoGS result. The sequence baselines and the SVM average about 0.68–0.86, and CAF-CNN is at 0.867. After dual-modality cross-attention aligns GPS, BDS, and GAL channels to the same tracking time axis, the fold means of PI-DMCANs lie between 0.847 and 0.988.

The ARR values in Table 6 place these drops against the within-scenario reference. The cross-scenario 17-fold mean of 0.932 remains above 0.90. A train-clean percentile does not by itself hold the false-alarm rate at 5% on a held-out scenario. Because the five baselines and PI-DMCANs share the same rule, the remaining gap combines threshold migration with missed detection on unseen attack families. CCC-PCNNs penalize one code-carrier relation on a one-dimensional sequence (Song et al. 2026). PI-DMCANs keeps that relation as one coordinate of a residual bank and adds the CCAF stream on the same channel. Correlator-image detectors such as CAF-CNN (Jin et al. 2024; Jin et al. 2026a) lose the matched-power and multi-constellation cases in which the heatmap stays close to the clean ridge. The tracking features and the residual gates recover those cases.

Once the spoofing scenario is held out under the cross-scenario protocol, the 17-fold Accuracy of CAF-CNN falls to 0.881, while PI-DMCANs remains at 0.932 (Table 5), about 5.1 percentage points above the strongest baseline. The ARR values in Table 6 place that result against the within-scenario reference (0.965), giving retention 0.966. The margin over the sequence models is largest on TEXBAT and MuCoGS; on OAKBAT several baselines are already near saturation. On matched-power and time-push cases such as ds7, every method falls, and missed detections dominate; a channel-wise score remains only weakly sensitive to receiver-common clock error, which is why $\tilde{r}_{\text{clk}}$ is kept out of $\mathbf{L}_{\text{cons}}$.

Within a single dataset, leaving one scenario out reduces PI-DMCANs from 0.965 to 0.928 (retention 0.962; Table 7), still about 3.5 percentage points above CAF-CNN at 0.893. Relative to the cross-dataset mean of 0.932, the within-dataset Accuracy changes only slightly, which indicates that most of the leave-one-scenario-out gain is not an artifact of training on a foreign dataset.

On the OAKBAT fold os6 (matched-power position spoofing under receiver motion), PI-DMCANs Accuracy falls to 0.872, below several baselines that remain near ceiling on that fold. The drop is not a loss of separability: detection probability remains 0.995 and AUC remains 0.997, whereas the false-alarm rate rises to 0.252 under the shared train-clean 95th-percentile threshold. Clean-channel scores on os6 sit higher than on the other OAKBAT folds, so the threshold estimated from os2–os5 intersects the elevated clean distribution. When os6 is held out, the within-dataset training pool contains static matched-power spoofing (os4) and dynamic overpowered spoofing (os5), but not the matched-power plus motion combination; motion-induced SQM and tracking residuals on the pre-injection clean segment are then read as weak spoofing by the residual-gated score. An oracle percentile taken on the os6 clean scores alone would restore Accuracy near 0.97, and the same phys schedule under the within-scenario split reaches about 0.975 on os6, which confirms that the fold is learnable once the operating point is matched to that scenario. The os6 result therefore illustrates the known limit

of a train-clean Neyman–Pearson rule under leave-one-scenario-out: the nominal 5% false-alarm rate need not hold after the score scale shifts, even when the ranking of clean and spoofed channels remains clear.

Relative to CCC-PCNNs, which penalize one code–carrier relation on a one-dimensional sequence, PI-DMCANs keeps that relation as one coordinate of a residual bank and adds the CCAF stream on the same channel. Relative to CAF-CNN (Jin et al. 2024; Jin et al. 2026a), the tracking stream and the residual gates recover cases in which the heatmap stays close to the clean ridge.

5.3 Implications, Limitations, and Future Work

For GNSS integrity tracking, the channel-wise NP decision is a basis for interference mitigation. A GNSS receiver can de-weight or exclude the flagged satellite before the navigation solution is formed, and it can do so on a single antenna, without an array, dual-antenna carrier-phase differences, or an inertial unit. Further work can adapt the monitor across front ends, couple it with sparse separation of authentic and counterfeit components and with position recovery, and close the loop from a channel-level alert to integrity de-weighting in the position domain.

6 Conclusions

PI-DMCANs performs channel-wise GNSS spoofing detection by aligning a short tracking sequence with a Mag/Q CCAF image and letting a residual bank set the two stream weights. Within-scenario accuracy at a Neyman–Pearson threshold was 0.965. Under within-dataset cross-scenario evaluation and cross-dataset evaluation, accuracy was 0.928 and 0.932 respectively, about 3.5 and 5.1 percentage points above the strongest baseline. Later work can adapt the monitor across front ends, link the channel flag to authentic and counterfeit separation, and pass that flag into the navigation filter.

Declarations

- Funding: This research did not receive funding.
- No moral or ethical issues exist.
- The authors declare that they have no competing interests.
- Consent for publication
- Data availability: TEXBAT and OAKBAT are publicly described benchmark datasets (Humphreys et al. 2012; Albright et al. 2020). MuCoGS is deposited in the Science Data Bank (<https://doi.org/10.57760/sciencedb.44739>) (Chen et al. 2026). Tracking-domain features were extracted with the open-source receiver FGI-GSRx (Islam et al. 2024). Processed detection scores and implementation code will be made available by the authors upon reasonable request.
- Materials availability: Not applicable
- Code availability: Not applicable
- Author contribution: Conceptualization, L.C. and W.M.; methodology, L.C. and Z.H.; software, L.C.; validation, L.C., Z.W. and W.M.; formal analysis, L.C.; investigation, L.C., Z.H. and Z.W.; resources, W.M.; data curation, L.C.; writing—original draft preparation, L.C.; writing—review and editing, L.C., G.O. and W.M.; visualization, L.C.; supervision, W.M.; project administration, W.M.; funding acquisition, W.M. All authors have read and agreed to the published version of the manuscript.

References

- Albright A, Powers S, Bonior J, Combs F (2020) A tool for furthering GNSS security research: The Oak Ridge spoofing and interference test battery (OAKBAT). In: 33rd international technical meeting of the satellite division of the institute of navigation (ION GNSS+ 2020). Institute of Navigation, Virtual, pp 3697–3712. <https://doi.org/10.33012/2020.17712>

- Bhuiyan MZH, Liaquat M, Islam S, Pääkkönen I, Saajasto M, Kaasalainen S (2026) Implementation and performance analysis of a Chi-Square Test based GNSS signal anomaly detection. *GPS Solut* 30:154. <https://doi.org/10.1007/s10291-026-02102-z>
- Borhani-Darian P, Li H, Wu P, Closas P (2024) Detecting GNSS spoofing using deep learning. *EURASIP J Adv Signal Process* 2024:14. <https://doi.org/10.1186/s13634-023-01103-1>
- Chen L, Ouyang X, Zeng F, Rui Z, Ming Y (2025) GNSS spoofing detection based on lightweight features and CGAN-ANN in unknown scenarios. *GPS Solut* 29:175. <https://doi.org/10.1007/s10291-025-01938-1>
- Chen L, Ou G, Lv Z, Hao Z, Mou W (2026) Raw multi-constellation GNSS spoofing dataset for civil navigation safety research (MuCoGS). *Science Data Bank*. <https://doi.org/10.57760/sciencedb.44739>
- Chen Q, Liu R, Yan Q, Xu Y, Liu Y, Huang X, Zhang Y (2025) Localization of multiple GNSS interference sources based on target detection in C/N distribution maps. *Remote Sens* 17:2627. <https://doi.org/10.3390/rs17152627>
- Chen Z, Li J, Li J, Zhu X, Li C (2022) GNSS multiparameter spoofing detection method based on support vector machine. *IEEE Sens J* 22:17864–17874. <https://doi.org/10.1109/JSEN.2022.3193388>
- Cho K, van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, Bengio Y (2014) Learning phrase representations using RNN encoder–decoder for statistical machine translation. In: 2014 conference on empirical methods in natural language processing (EMNLP). Association for Computational Linguistics, Doha, pp 1724–1734. <https://doi.org/10.3115/v1/D14-1179>
- Cui B, Jiang X, Wang J, Li P, Ge M, Schuh H (2023) A new large-area hierarchical PPP-RTK service strategy. *GPS Solut* 27:132. <https://doi.org/10.1007/s10291-023-01476-8>
- Feng F, Chen Z, Chen L, Zhu X (2024) Rapid detection of GNSS time synchronization attacks via the enhanced code and carrier Doppler consistency test. *GPS Solut* 28:38. <https://doi.org/10.1007/s10291-023-01583-6>
- Gao Y, Lv Z, Zhang L (2020) Asynchronous lift-off spoofing on satellite navigation receivers in the signal tracking stage. *IEEE Sens J* 20:8604–8613. <https://doi.org/10.1109/JSEN.2020.2984525>
- Graves A, Schmidhuber J (2005) Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Netw* 18:602–610. <https://doi.org/10.1016/j.neunet.2005.06.042>
- Hao Z, Liu X, Zhao X, Li Z, Lu Z (2026) GNSS spoofing interference detection based on transformer multi-parameter weighted fusion. *Syst Eng Electron* 48:2434–2447. <https://doi.org/10.12305/j.issn.1001-506X.2026.07.27>
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9:1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Huang G, Sun Y, Liu Z, Sedra D, Weinberger KQ (2016) Deep networks with stochastic depth. In: Leibe B, Matas J, Sebe N, Welling M (eds) *Computer vision – ECCV 2016*. Springer, Cham, pp 646–661. https://doi.org/10.1007/978-3-319-46493-0_39
- Humphreys TE, Bhatti JA, Shepard DP, Wesson KD (2012) The Texas spoofing test battery: Toward a standard for evaluating GPS signal authentication techniques. In: 25th international technical meeting of the satellite division of the institute of navigation (ION GNSS 2012). Institute of Navigation, Nashville, TN, USA, pp 3569–3583.
- Iqbal A, Aman MN, Sikdar B (2024a) Machine and representation learning-based GNSS spoofing detectors utilizing feature set from generic GNSS receivers. *IEEE Trans Consum Electron* 70:574–589. <https://doi.org/10.1109/TCE.2023.3346287>
- Iqbal A, Aman MN, Sikdar B (2024b) A deep learning based induced GNSS spoof detection framework using uniform experimental data. *IEEE Trans Mach Learn Commun Netw* 2:465–478. <https://doi.org/10.1109/TMLCN.2024.3386649>
- Islam S, Bhuiyan MZH, Liaquat M, Pääkkönen I, Kaasalainen S (2024) An open GNSS spoofing data repository: Characterization and impact analysis with FGI-GSRx open-source software-defined receiver. *GPS Solut* 28:193. <https://doi.org/10.1007/s10291-024-01719-2>
- Ji N, Rao Y, Wang X, Zou D (2024) Advanced GNSS spoofing detection: Aggregated correlation residue likelihood analysis. *Remote Sens* 16:2868. <https://doi.org/10.3390/rs16152868>
- Jin R, Cui X, Yan J, Xiong H, Yang H, Gu M, Zhen W (2024) Small delay GNSS forwarding spoofing detection in a multipath environment based on convolutional neural network. *IEEE Sens J* 24:24070–24085. <https://doi.org/10.1109/JSEN.2024.3407811>
- Jin R, Yan J, Cui X, Yang H, Zhen W, Gu M, Ji G, Chen L, Li H (2025) A spoofing detection and direction-finding approach for global navigation satellite system signals using off-the-shelf anti-jamming antennas. *Remote Sens* 17:864. <https://doi.org/10.3390/rs17050864>
- Jin X, Zhang X, Wang W, Qi Z, Zheng S (2026a) SpoofConvNet-based enhanced signal quality monitoring for GNSS spoofing detection in consumer electronics. *IEEE Trans Consum Electron* 72:1279–1291. <https://doi.org/10.1109/TCE.2025.3649205>
- Jin X, Zhang X, Zheng S (2026b) Indirect Kalman filtering for robust GNSS spoofing detection in signal quality monitoring. *Signal Process* 239:110321. <https://doi.org/10.1016/j.sigpro.2025.110321>
- LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86:2278–2324. <https://doi.org/10.1109/5.726791>
- Li J, Zhu X, Ouyang M, Li W, Chen Z, Dai Z (2020) Research on multi-peak detection of small delay spoofing signal. *IEEE Access* 8:151239–151251. <https://doi.org/10.1109/ACCESS.2020.3016971>

- Li J, Chen Z, Yuan X, Xie T, Xu Y, Zheng Z, Zhu X (2025) A real-time GNSS time spoofing detection framework based on feature processing. *GPS Solut* 29:45. <https://doi.org/10.1007/s10291-024-01802-8>
- Li S, Tang X, Lin H, Wang F (2025) GNSS spoofing detection based on frequency domain processing. *Meas* 242:115872. <https://doi.org/10.1016/j.measurement.2024.115872>
- Lin T-Y, Goyal P, Girshick R, He K, Dollár P (2017) Focal loss for dense object detection. In: *IEEE international conference on computer vision (ICCV)*. IEEE, Venice, pp 2999–3007. <https://doi.org/10.1109/ICCV.2017.324>
- Liu R, Yang Z, Chen Q, Liao G, Zhu Q (2023) Localization of GNSS spoofing interference source based on a moving array antenna. *Remote Sens* 15:5497. <https://doi.org/10.3390/rs15235497>
- Mao W, Ren J, Ni S (2025) Fast GNSS spoofing detection based on LSTM-detect model. *GPS Solut* 29:57. <https://doi.org/10.1007/s10291-025-01819-7>
- Marata L, Bhuiyan MZH, Lohan ES (2025) On cross-testing datasets for RF-fingerprinting based deep-learning GNSS spoofing detection. In: *2025 IEEE international workshop on signal processing and artificial intelligence for wireless communications (SPAWC)*. IEEE, Surrey, pp 1–5. <https://doi.org/10.1109/SPAWC66079.2025.11143386>
- Pervysheva Y, Käppi J, Syrjärinne J, Nurmi J, Lohan ES (2024) Multi-receiver ensemble machine-learning techniques for GNSS spoofing detection with generalisability approach. In: *2024 32nd telecommunications forum (TELFOR)*. IEEE, Belgrade, pp 1–4. <https://doi.org/10.1109/TELFOR63250.2024.10819159>
- Psiaki ML, Humphreys TE (2016) GNSS spoofing and detection. *Proc IEEE* 104:1258–1270. <https://doi.org/10.1109/JPROC.2016.2525516>
- Radoš K, Brkić M, Begušić D (2024) Recent advances on jamming and spoofing detection in GNSS. *Sensors* 24:4210. <https://doi.org/10.3390/s24134210>
- Semanjski S, Semanjski I, De Wilde W, Muls A (2019) Use and validation of supervised machine learning approach for detection of GNSS signal spoofing. In: *2019 international conference on localization and GNSS (ICL-GNSS)*. IEEE, Nuremberg, pp 1–6. <https://doi.org/10.1109/ICL-GNSS.2019.8752775>
- Shang X, Sun F, Wang D, Xiao K, Dou S, Lu X, Sun J (2023) GNSS spoofing detection based on multicorrelator distortion monitoring. *GPS Solut* 27:94. <https://doi.org/10.1007/s10291-023-01438-0>
- Song Y, Xiao W, Ye X, Yang X, Liu W, Sun G (2026) Generalized GNSS spoofing detection in unseen scenarios using physics-constrained neural networks with code-carrier consistency. *Satell Navig* 7:15. <https://doi.org/10.1186/s43020-026-00199-8>
- Sun C, Cheong JW, Dempster AG, Demicheli L, Cetin E, Zhao H, Feng W (2018a) Moving variance-based signal quality monitoring method for spoofing detection. *GPS Solut* 22:83. <https://doi.org/10.1007/s10291-018-0745-7>
- Sun C, Cheong JW, Dempster AG, Zhao H, Feng W (2018b) GNSS spoofing detection by means of signal quality monitoring (SQM) metric combinations. *IEEE Access* 6:66428–66441. <https://doi.org/10.1109/ACCESS.2018.2875948>
- Sun C, Cheong JW, Dempster AG, Zhao H, Bai L, Feng W (2021) Robust spoofing detection for GNSS instrumentation using Q-channel signal quality monitoring metric. *IEEE Trans Instrum Meas* 70:8504115. <https://doi.org/10.1109/TIM.2021.3102753>
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: *Advances in neural information processing systems*. Curran Associates, Red Hook, pp 5998–6008. <https://doi.org/10.5555/3295222.3295349>
- Vemuri SS, Nurmi J, Bhuiyan MZH, Islam S, Lohan ES, Torres-Sospedra J (2026) A deep reinforcement learning framework for GNSS spoofing detection and classification. In: *Proceedings of the 2026 international technical meeting of the Institute of Navigation*. Institute of Navigation, Anaheim, pp 414–427. <https://doi.org/10.33012/2026.20554>
- Wang L, She J, Cui B, Wang Z, Yang W, Wang Y (2025) Mitigating integrity risk in SBAS positioning using enhanced IGG III robust estimation. *Remote Sens* 17:3067. <https://doi.org/10.3390/rs17173067>
- Wang X, Yang J, Peng Z, Wang S, Huang M (2024) Hilbert signal envelope-based multi-features methods for GNSS spoofing detection. *Comput Secur* 144:103959. <https://doi.org/10.1016/j.cose.2024.103959>
- Wesson KD, Gross JN, Humphreys TE, Evans BL (2018) GNSS signal authentication via power and distortion monitoring. *IEEE Trans Aerosp Electron Syst* 54:739–754. <https://doi.org/10.1109/TAES.2017.2765258>
- Wu M, Wang L, Xie W, Yue F, Cui B (2024) Performance evaluation and application field analysis of precise point positioning based on different real-time augmentation information. *Remote Sens* 16:1349. <https://doi.org/10.3390/rs16081349>
- Yang B, Dong C, Gao B, Liu Y, Cui W, Gao F (2022) Research on GNSS interference recognition based on ROI of correlation peaks. *Int J Satell Commun Netw* 40:349–363. <https://doi.org/10.1002/sat.1444>
- Zhang C, Wang D, Wu J (2024) Two-dimensional directions determination for GNSS spoofing source based on MEMS-based dual-GNSS/INS integration. *Remote Sens* 16:4568. <https://doi.org/10.3390/rs16234568>
- Zhu X, Lu Z, Hua T, Yang F, Tu G, Chen X (2022) A novel GPS meaconing spoofing detection technique based on improved ratio combined with carrier-to-noise moving variance. *Electronics* 11:738. <https://doi.org/10.3390/electronics11050738>