Deep-Unfolded Beamforming for GNSS Jamming and Spoofing Mitigation

Janghyuk Youn
Dept. of AI Convergence Network
Ajou University
Suwon, Republic of Korea
jhyoun@ajou.ac.kr

Bang Chul Jung
Dept. of Electrical and Computer Engineering
Ajou University
Suwon, Republic of Korea
bcjung@ajou.ac.kr

***Abstract*—Global navigation satellite system (GNSS) receivers remain vulnerable to diverse jamming and spoofing attacks, while existing countermeasures are often evaluated under inconsistent signal conditions and predominantly using radio-frequency metrics. In this paper, we propose a unified evaluation platform and a deep-unfolded beamformer for GNSS jamming and spoofing mitigation. The platform synthesizes array signals under five interference types and evaluates arbitrary beamformers through a common metric chain spanning output signal-to-interference-plus-noise ratio (SINR), pseudorange precision, and spoofing capture, thereby assessing radio-frequency suppression and navigation integrity on the same signal realization. The proposed beamformer unrolls the minimum variance distortionless response (MVDR) solution into a fixed number of layers and learns layer-wise step sizes and diagonal loading from covariance statistics without requiring labeled weight vectors. Simulation results show that output SINR alone may not reliably predict spoofing capture. Under the same iteration budget as the conjugate-gradient recursion, the learned schedule improves output SINR by up to 6.5 dB, with larger gains for larger arrays under the tightest iteration constraint.**

***Index Terms*—GNSS, anti-jamming, anti-spoofing, controlled reception pattern antenna, adaptive beamforming, deep unfolding, output SINR**

I. Introduction

The received power of a global navigation satellite system (GNSS) signal at the Earth's surface is 27.6 dB below the thermal noise floor of a 10.23 MHz front end with a 3 dB noise figure. Consequently a jammer only tens of decibels above that floor denies service over a wide area, and a spoofer imposes a false position without being noticed at all, which now puts autonomous vehicles, unmanned aerial systems and critical infrastructure at risk [1]. Single-antenna countermeasures excise interference that is sparse in time or in frequency, however wideband Gaussian jamming is sparse in neither. Array processing with a controlled reception pattern antenna (CRPA) is therefore the only general defense, since such interference remains sparse in the spatial domain [2].

The same spatial degrees of freedom answer both threats. Against spoofing the receiver estimates the direction of arrival (DoA) of every pseudo-random noise signal and declares an attack when several of them share one direction [3], after which a minimum mean squared error beamformer removes the counterfeit signal [4]. That line is limited by the DoA resolution, which degrades for closely spaced sources and short observation windows and calls for a hierarchical subspace search [5]. Against wideband jamming no source is resolved individually, and constrained least mean square (LMS) adaptation, power inversion (PI) and the minimum-variance distortionless response (MVDR) beamformer work from the sample covariance, the last in closed form through its inverse. Learning-based processing is more recent. In [2], a long short-term memory network estimates the aggregate interference when the interferers outnumber the antennas, which improves on classical processing but replaces an interpretable estimator with a black-box network whose parameter count and generalization behavior are difficult to characterize.

Iterative alternatives to the matrix inverse are well established. Reduced-rank processing projects the weights onto the Krylov subspace spanned by the covariance and the constraint vector, removing both the eigenvalue decomposition and the inversion. The multistage Wiener filter [6] and the conjugate gradient (CG) recursion are equivalent in this sense [7], and both fix the step rule analytically.

Two gaps remain. First, the two threats are met by separate techniques reported under mutually incompatible signal conditions, and the anti-jamming branch is almost always judged by the output signal-to-interference-plus-noise ratio (SINR) alone, so a technique that leaves the code-tracking loop captured by a spoofer is indistinguishable from one that does not. Second, deep unfolding [8] has reached adaptive beamforming [9] and learned steps are known to reproduce Chebyshev acceleration [10], yet to the best of our knowledge what a learned schedule buys over an analytic one at an equal iteration budget has not been quantified for GNSS array anti-jamming.

We address these gaps with a unified platform that evaluates beamformers on identical signal realizations using common radio-frequency and navigation metrics, and a deep-unfolded beamformer that learns layer-wise step sizes and diagonal loading from covariance statistics. The proposed design targets output-SINR performance under a fixed iteration budget rather than asymptotic convergence.

II. Evaluation Platform and Signal Model

The platform composes the array signal, runs an arbitrary set of beamformers on it, and reports the metric chain of

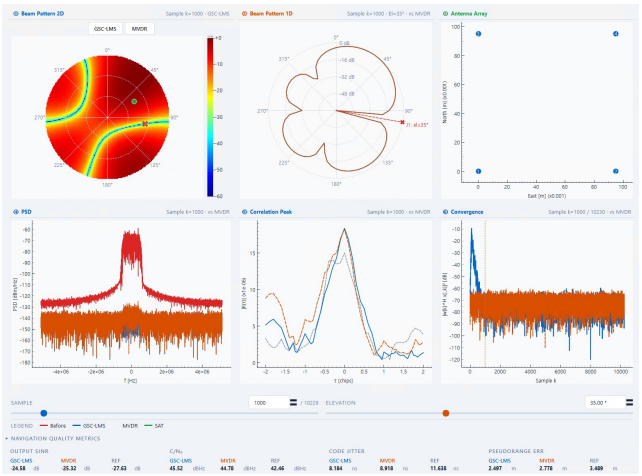

Fig. 1. Interface of the evaluation platform, showing the beam patterns, the array geometry, the spectrum, the correlation peak, and the navigation quality metrics.

Section II-B, as shown in Fig. 1. A technique joins the comparison by providing a single weight-computation routine.

A. Signal and Interference Model

We consider an M-element CRPA receiving one desired satellite and K interferers. The sampled array signal $\mathbf{x}[n] \in \mathbb{C}^M$ is modeled as

$$\mathbf{x}[n] = \sqrt{P_s}\,\mathbf{a}_s\,s[n] + \sum_{k=1}^{K} \sqrt{P_k}\,\mathbf{a}_k\,\tilde{j}_k[n] + \mathbf{n}[n], \qquad (1)$$

where $s[n]$ and $\tilde{j}_k[n]$ are the unit-power baseband waveforms of the satellite and of the kth interferer, P_s is the received satellite power specified by the interface control document, P_k is the received power of the kth interferer, and $\mathbf{n}[n] \sim \mathcal{CN}(\mathbf{0}, \sigma_n^2 \mathbf{I}_M)$ is thermal noise with $\sigma_n^2 = k_B T_0 B F$, in which B is the front-end bandwidth and F the noise figure. The steering vector $\mathbf{a}(\theta, \phi)$ has unit-modulus entries $\exp(j\psi_m)$ with $\psi_m = \frac{2\pi}{\lambda}(x_m \cos\theta \sin\phi + y_m \cos\theta \cos\phi + z_m \sin\theta)$ in the navigation convention, evaluated at the satellite and interferer directions to give $\mathbf{a}_s$ and $\mathbf{a}_k$. Five interference types are generated, namely continuous wave, band-limited Gaussian, pulsed, chirp, and a spoofer replaying the desired code with a delay and a Doppler offset. The jammer-to-signal ratio (JSR) is referenced to the satellite power, so a jammer becomes a threat only once it exceeds the 27.6 dB gap to the noise floor.

Every beamformer observes the same realization and produces $\mathbf{w} \in \mathbb{C}^M$ with output $y[n] = \mathbf{w}^H \mathbf{x}[n]$. The weights are computed from the first N snapshots and scored on a disjoint block, which keeps the estimation and measurement windows separable.

B. Metric Chain to the Navigation Observable

The platform evaluates each weight vector along a single chain. The output SINR is

$$\mathrm{SINR} = \frac{P_s |\mathbf{w}^H \mathbf{a}_s|^2}{\mathbf{w}^H \mathbf{R}_{j+n} \mathbf{w}}, \qquad (2)$$

where $\mathbf{R}_{j+n}$ is the measured covariance of the scoring block with the known desired-signal contribution $P_s \mathbf{a}_s \mathbf{a}_s^H$ removed, a spoofer being retained in it as interference. This construction enters the score alone and never the weights. The carrier-to-noise density follows as $C/N_0 = \mathrm{SINR} + 10\log_{10} B$, and the delay-locked loop jitter for a non-coherent early-minus-late discriminator with loop bandwidth B_L, correlator spacing d and coherent integration time T is [11]

$$\sigma_\tau = \sqrt{\frac{B_L d}{2\,(C/N_0)} \left(1 + \frac{2}{(2-d)\,T\,(C/N_0)}\right)} \qquad (3)$$

in chips, which maps to a pseudorange precision of $c\,T_c\,\sigma_\tau$. The second factor is the squaring loss, which is not negligible at the low C/N_0 that strong jamming produces.

Precision alone does not capture spoofing. We therefore locate the dominant correlation peak of the post-beamformer output against the local replica and declare a *spoofer capture* when it lies within 0.2 chip of the spoofer code delay, which indicates that the tracking loop has locked onto the counterfeit signal even when the output SINR is unaffected.

III. Deep-Unfolded Beamforming Technique

A. Unrolled MVDR Iteration

The MVDR weight vector solves $\mathbf{R}\mathbf{w} \propto \mathbf{a}_s$ for the sample covariance $\mathbf{R} = N^{-1} \sum_{n=1}^{N} \mathbf{x}[n]\mathbf{x}^H[n]$ of the estimation block, so the matrix inverse can be replaced by a fixed number of Richardson steps,

$$\mathbf{w}_{l+1} = \mathbf{w}_l + \gamma_l\,(\mathbf{a}_s - \mathbf{R}_\alpha \mathbf{w}_l), \quad l = 0, \ldots, L-1, \qquad (4)$$

with $\mathbf{w}_0 = \mathbf{a}_s/M$ and $\mathbf{R}_\alpha = \tilde{\mathbf{R}} + \alpha \mathbf{I}_M$, where $\tilde{\mathbf{R}} = M\mathbf{R}/\mathrm{tr}(\mathbf{R})$ is the trace-normalized sample covariance. The normalization removes the absolute power scale, so a single parameter set applies across the whole JSR range. A final projection $\mathbf{w} \leftarrow \mathbf{w}/(\mathbf{a}_s^H \mathbf{w})$ enforces the distortionless constraint.

Unrolling (4) to L layers yields a feed-forward network of fixed depth whose reachable set is the Krylov subspace $\mathrm{span}\{\mathbf{a}_s, \mathbf{R}_\alpha \mathbf{a}_s, \ldots, \mathbf{R}_\alpha^L \mathbf{a}_s\}$. A budget of L products fixes that set, so every method spending it differs only in where inside the set it lands. The CG recursion lands on the point that minimizes the quadratic $\frac{1}{2}\mathbf{w}^H \mathbf{R}_\alpha \mathbf{w} - \Re\{\mathbf{w}^H \mathbf{a}_s\}$ [7], whereas the beamforming objective is the output power under the distortionless constraint. The two minimizers coincide only when the subspace is complete, so room is left whenever the budget is short of that. The fixed choice $\gamma_l = 1/\lambda_{\max}(\mathbf{R})$ of the conventional unrolling is weaker still, since its convergence factor degrades with the condition number, which grows with the JSR, and the Chebyshev schedule that learned steps recover [10] needs spectral edges a receiver does not have.

B. Learned Schedule

We call the resulting beamformer DUNA, for deep-unfolded nulling for anti-jamming, and let its steps and loading depend on the observed covariance through four statistics, the normalized power sums $p_k = \mathrm{tr}(\tilde{\mathbf{R}}^k)/M$ for $k = 2, 3, 4$ and $\log\det\tilde{\mathbf{R}}$. The power sums are what Newton's identities

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Band, chip rate	GPS L1 C/A, 1.023 Mcps
Sampling rate, noise figure	10.23 MHz, 3 dB
Satellite power (IS-GPS-200)	-128.5 dBm
Array, interferers	M-element URA (0.5λ), $K \leq 3$
JSR, window N, budget L	30–80 dB, 10,230, 1–8
Diagonal loading (MVDR, PI, CG)	-40 dB rel. mean element power
LMS step size	0.3
Monte Carlo runs	100 per point (50 in Fig. 3)

need to write the characteristic polynomial, and so in principle $\mathbf{R}^{-1}$ itself. They weight the eigenvalues by their kth power, however, so the small ones vanish and the set saturates once the jammers clear the noise floor. The log-determinant sums the log-eigenvalues instead and does not saturate over the JSR range considered here.

A network g_θ with one hidden layer of sixteen $\tanh$ units maps the standardized statistics to a relative correction of the steps and to the loading,

$$\gamma_l = \bar{\gamma}_l \left(1 + [g_\theta]_l\right), \tag{5}$$

$$\alpha = \mathrm{softplus}([g_\theta]_{L+1}), \tag{6}$$

which is $98 + 18L$ parameters, or 134 at the budget of Fig. 4. The correction is multiplicative because the learned steps span two orders of magnitude across layers, which an additive head cannot cover under one learning rate.

Training minimizes the MVDR objective itself,

$$\mathcal{L}(\theta) = \mathbb{E}\left[\log \mathbf{w}^H \tilde{\mathbf{R}}_{\mathrm{val}} \mathbf{w}\right], \tag{7}$$

evaluated on a covariance from a block disjoint from the one used to compute $\mathbf{w}$. No reference weight vector is required, so the scheme never asks for a quantity that a receiver cannot measure. The logarithm makes the batch loss an average over decibels, so the scenarios with the largest absolute output power do not dominate.

IV. Simulation Results

A. Setup

Table I lists the parameters. The sampling rate follows the 1–6 MHz interference bandwidth that is swept rather than the 2.046 MHz main lobe of the C/A code. All weight vectors are normalized to unit norm before scoring so that the output noise variance is identical. The adaptive baseline is Frost's constrained LMS, which projects every update through $\mathbf{P} = \mathbf{I} - \mathbf{a}_s \mathbf{a}_s^H / M$, and PI is $\mathbf{w} = \mathbf{R}^{-1}\mathbf{e}_r / (\mathbf{e}_r^H \mathbf{R}^{-1} \mathbf{e}_r)$ with $\mathbf{e}_r$ the first array element. Each baseline receives one hyperparameter selected by a grid search on the training distribution, in the same way that the proposed scheme receives one parameter set for that distribution. MVDR therefore acts as an implementable upper bound rather than a genie.

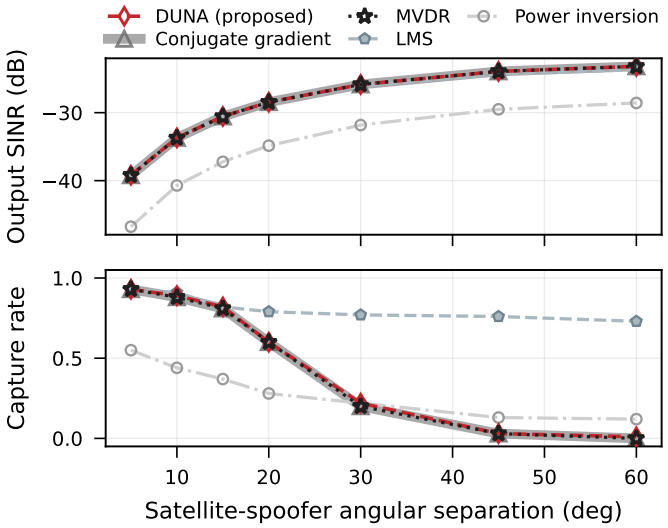

Fig. 2. Output SINR and spoofer capture rate versus satellite-spoofer angular separation, when the spoofer is the only interferer, its JSR is 50 dB, the code delay is 0.5 chip and $N = 10{,}230$.

B. Platform: the Ranking Depends on the Scenario

The platform first shows that no single technique dominates. Constrained LMS stays within 1 dB of MVDR for continuous-wave, chirp and wideband interference, however it loses 14.1 dB under pulsed interference because the instantaneous power during the pulse drives the effective step size past the stability bound. PI, which trails LMS by about 3 dB on the other three types, overtakes it by 10.6 dB there, so the ranking inverts with the interference type.

Fig. 2 shows the second finding. At a separation of $60°$, LMS and MVDR differ by 0.49 dB in output SINR, yet LMS is captured by the spoofer in 0.73 of the trials, with a 95% interval of $[0.64, 0.82]$, and MVDR in none of them. The two disagree on 73 of the 100 realizations and never in the opposite direction. PI inverts the ordering outright, since it is 4.9 dB worse than LMS in output SINR while being captured in only 0.12 $[0.06, 0.18]$ of the trials. The mechanism is that LMS only has to reduce output power, which a spoofer carrying the satellite's own code already does, so the adaptation has no reason to remove it. MVDR, CG and DUNA instead see one more direction in the covariance and null it, and are captured in at most 0.01 of the trials. This disagreement is invisible to any evaluation that stops at the radio-frequency metric.

C. Equal-Budget Comparison

Fig. 3 reports the loss with respect to the closed form as a function of the budget, for three array sizes. At one iteration the learned schedule leads the CG recursion by 2.38, 4.33 and 6.46 dB for $M = 4$, 8 and 16, and at two iterations by 2.08, 0.93 and 1.58 dB. The lead therefore grows with the array at the tightest budget, which is where a fixed pipeline operates.

The two meet once the budget reaches the number of eigenvalue clusters, of which the covariance carries $K + 1$, one per interferer plus noise. A polynomial iteration needs as many steps as there are clusters rather than as many as there

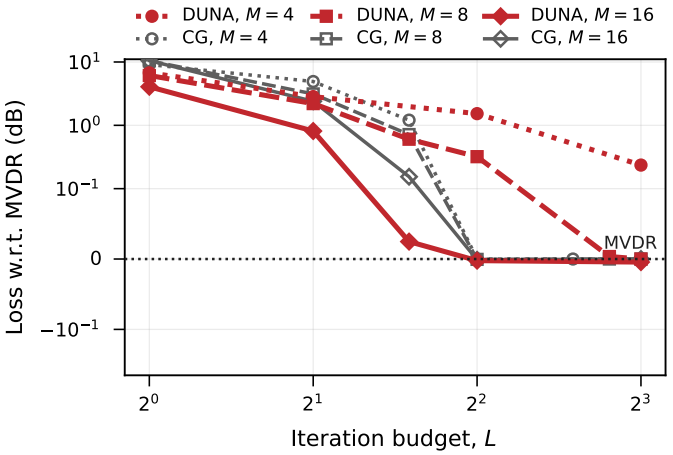

Fig. 3. Loss with respect to the closed-form MVDR versus iteration budget, when $M \in \{4, 8, 16\}$, $K \in \{1, 2, 3\}$, the JSR is 50 dB and $N = 10{,}230$.

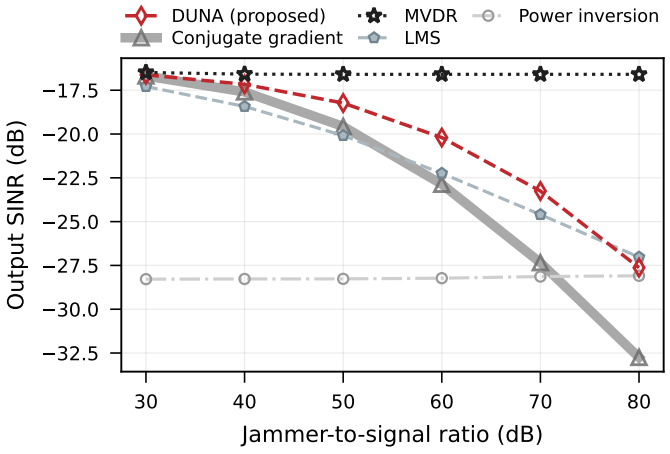

Fig. 4. Output SINR versus jammer-to-signal ratio at a fixed budget of $L = 2$ iterations, when $M = 16$, $K \in \{1, 2, 3\}$ and $N = 10{,}230$.

are antennas, and beyond that point CG has terminated at the exact solution with nothing left for a schedule to learn. At $M = 4$ and $M = 8$ it then passes the learned schedule by 1.53 and 0.32 dB at $L = 4$. We report the crossing rather than truncating the axis at it, since the regime the proposal addresses is the one below it.

D. Robustness at a Fixed Budget

Fig. 4 holds the budget at two iterations and sweeps the jamming level for $M = 16$. The margin over CG widens monotonically from 0.06 dB at a JSR of 30 dB to 5.12 dB at 80 dB, because the covariance grows more ill-conditioned and an analytically fixed step rule has no way to react to it. Against the classical baselines the learned schedule leads constrained LMS by 0.67 to 2.04 dB up to a JSR of 70 dB and power inversion by 11.66 dB at the weakest jamming.

Two iterations are not free of cost. The proposal stays within 0.14 dB of the closed form at 30 dB but falls 11.03 dB behind it at 80 dB, where LMS also overtakes it by 0.60 dB. The claim is therefore not that two layers replace the matrix inverse, but that once the budget is fixed the schedule is worth several decibels, and most where jamming is strongest.

V. Conclusion

In this paper, we have proposed a unified evaluation platform for GNSS anti-jamming and anti-spoofing together with a deep-unfolded beamformer whose step sizes and diagonal loading are learned from covariance statistics without labeled weights. The platform has shown that the ranking of techniques depends on the interference type and that the output SINR and the spoofer capture rate disagree under spoofing. At an equal iteration budget the learned schedule has gained up to 6.46 dB over the conjugate gradient recursion, with the gain growing with the array size at the tightest budget and with the jamming strength at a fixed one. As future work, we will extend the scheme to space-time adaptive processing and evaluate it on recorded live-sky data.

Acknowledgment

This work was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea Government (MSIT) (No.RS-2024-00396992, No.RS-2026-25525599); and in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.RS-2026-25489627).

References

[1] B. C. Jung, "Beyond GNSS: Cellular/NTN positioning, LEO-PNT, and AI for resilient wireless localization," *IEEE Veh. Technol. Mag.*, vol. 21, no. 2, pp. 5–11, Jun. 2026.

[2] Y. Xie, L. Chen, K. Zhang, X. Huang, J. Peng, and S. Ni, "Deep neural network based anti-jamming processing in the GNSS array receiver," *GPS Solut.*, vol. 29, no. 2, 2025.

[3] Y.-S. Lee, J. S. Yeom, J. H. Noh, S. J. Lee, and B. C. Jung, "A novel GNSS spoofing detection technique with array antenna-based multi-PRN diversity," *J. Position. Navig. Timing*, vol. 10, no. 3, pp. 169–177, Sep. 2021.

[4] Y.-S. Lee, J. S. Yeom, and B. C. Jung, "A novel array antenna-based GNSS spoofing detection and mitigation technique," in *Proc. IEEE Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2023, pp. 489–492.

[5] M. Oh, Y.-S. Lee, J.-B. Lim, C.-O. Kang, and B. C. Jung, "Hierarchical MUSIC technique for direction finding of multiple signal sources," in *Proc. IEEE Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2024, pp. 1082–1083.

[6] J. S. Goldstein, I. S. Reed, and L. L. Scharf, "A multistage representation of the Wiener filter based on orthogonal projections," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2943–2959, Nov. 1998.

[7] M. E. Weippert, J. D. Hiemstra, J. S. Goldstein, and M. D. Zoltowski, "Insights from the relationship between the multistage Wiener filter and the method of conjugate gradients," in *Proc. IEEE Sensor Array Multichannel Signal Process. Workshop*, Aug. 2002, pp. 388–392.

[8] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, Mar. 2021.

[9] R. Janani and R. Fatemi Mofrad, "Deep unfolding model-based for hybrid robust wide band adaptive beamforming," *IET Microw. Antennas Propag.*, vol. 18, no. 7, pp. 480–493, 2024.

[10] S. Takabe and T. Wadayama, "Convergence acceleration via Chebyshev step: Plausible interpretation of deep-unfolded gradient descent," *IEICE Trans. Fundamentals*, vol. E105.A, no. 8, pp. 1110–1120, Aug. 2022.

[11] E. D. Kaplan and C. J. Hegarty, Eds., *Understanding GPS/GNSS: Principles and Applications*, 3rd ed. Norwood, MA, USA: Artech House, 2017.