

Sturdy RF Direction Finding in WSNs Under Multipath and Non-Line-of-Sight Conditions Using Machine Learning

Yaseen Naser Jurn^{1*}, Abeer Tariq Maolood², Ekhlas Khalaf Gbashi², Sawsen Abdulhadi Mahmood³

¹College of Engineering, University of Information Technology and Communications, Baghdad, Iraq

²College of Computer Science, University of Technology, Baghdad, Iraq

³Department of Computer Science, College of Education, Mustansiriyah University, Baghdad, Iraq

DOI: <https://doi.org/10.51584/IJRIAS.2026.11080152>

Received: 02 September 2026; Accepted: 07 September 2026; Published: 19 September 2026

ABSTRACT

Accurate radio-frequency (RF) direction finding (DF) is a cornerstone capability for localization, tracking, and situational awareness in wireless sensor networks (WSNs). Classical subspace and beamforming-based angle-of-arrival (AoA) estimators such as MUSIC and ESPRIT degrade sharply once the propagation channel departs from the idealized single-path, line-of-sight (LOS) assumption on which they are built. In dense, cluttered, or urban WSN deployments, multipath fading and non-line-of-sight (NLOS) propagation routinely corrupt the array manifold, biasing angle estimates by tens of degrees and destabilizing downstream localization. This paper presents a sturdy, learning-based direction-finding framework that jointly (i) detects and classifies the propagation condition (LOS, weak-NLOS, strong-NLOS) from low-level array statistics, and (ii) regresses the angle of arrival using a hybrid convolutional–recurrent network trained to be invariant to multipath-induced covariance distortion. The proposed pipeline ingests the sample covariance matrix and eigen-spectrum of a uniform circular sensor array (UCA), which offers unambiguous 360° azimuthal coverage well suited to WSN nodes with no prior knowledge of source half-plane, applies a lightweight convolutional feature extractor to capture spatial correlation structure, and uses a bidirectional long short-term memory (Bi-LSTM) stage together with an attention layer to resolve angular ambiguities introduced by correlated multipath returns. An auxiliary NLOS-classification head regularizes training and enables condition-aware error correction at inference time. Extensive simulations across a Saleh–Valenzuela multipath channel model, varying signal-to-noise ratio (SNR), number of multipath taps, K-factor, and array geometry, show that the proposed model reduces root-mean-square angular error by 46–68% relative to MUSIC, ESPRIT, and a plain feed-forward DoA regressor, while sustaining sub-3° RMSE down to 0 dB SNR and under severe NLOS conditions. We further demonstrate resilience to sensor-count reduction and model compactness suitable for resource-constrained WSN nodes.

Index Terms—Direction finding, angle of arrival, wireless sensor networks, multipath, non-line-of-sight, deep learning, CNN-LSTM, robust localization.

INTRODUCTION

Wireless sensor networks (WSNs) increasingly rely on RF direction finding (DF) — i.e., angle-of-arrival (AoA) estimation — as a core primitive for source localization, target tracking, beam steering, and topology-aware routing. Applications ranging from search-and-rescue, industrial asset tracking, smart-building occupancy sensing, to battlefield situational awareness depend on WSN nodes being able to bear a transmitting source with sufficient angular resolution using compact, low-cost antenna arrays.

Classical DF algorithms — delay-and-sum beamforming, Capon’s minimum-variance distortionless response (MVDR), and subspace methods such as MUSIC [1] and ESPRIT [2] — achieve high resolution under the idealized assumption of a small number of discrete, uncorrelated, far-field plane waves impinging on a well-calibrated array in an otherwise line-of-sight (LOS) channel. This assumption is routinely violated in realistic WSN deployment environments: indoor offices, urban canyons, forested terrain, and industrial floors introduce

rich multipath scattering, specular reflections, and frequent non-line-of-sight (NLOS) blockage of the direct path. Under these conditions, (i) the received signal covariance matrix becomes rank-deficient or highly correlated due to coherent multipath, collapsing the noise subspace on which MUSIC/ESPRIT depend, and (ii) the apparent AoA of the dominant received cluster is biased away from the true bearing, sometimes by tens of degrees, because the strongest path is a reflection rather than the direct path. The net effect is a DF subsystem that is accurate in controlled testbeds but whose sturdy performance collapses in the field — precisely where WSNs are most needed.

Machine learning (ML) offers an attractive route to sturdiness because it can learn the mapping from corrupted, multipath-contaminated array statistics directly to bearing angle, without relying on a clean-signal subspace assumption. Recent work has explored deep neural networks, convolutional architectures, and residual/attention models for DoA estimation under noise and limited snapshots [5]–[9], and several efforts specifically target NLOS-corrupted AoA estimation [3, 10] and NLOS identification for ranging/positioning [11]–[13]. However, much of this literature targets either (a) generic low-SNR/limited-snapshot regimes without an explicit multipath/NLOS channel model, or (b) NLOS classification for range-based positioning rather than direct AoA regression, or (c) single large arrays unsuitable for resource-constrained WSN nodes. There remains a gap for a DF framework that: explicitly models WSN-relevant multipath channels (e.g., Saleh–Valenzuela clustered-arrival statistics), jointly performs NLOS-condition awareness and angle regression in one lightweight pipeline, and is evaluated for sturdiness across the full space of SNR, K-factor, and multipath density that a deployed WSN node will encounter.

Contributions. This paper makes the following contributions:

- We formulate RF direction finding in WSNs as a joint condition-classification and angle-regression learning problem, and design a compact hybrid CNN–BiLSTM–attention architecture (RD-Net) that consumes covariance/eigen-statistics rather than raw IQ samples, keeping the input dimensionality and model size suitable for embedded WSN hardware.
- We introduce an NLOS-aware auxiliary loss and a condition-conditioned bias-correction head that explicitly compensates for the systematic AoA bias introduced by dominant reflected paths.
- We build a WSN-realistic multipath/NLOS simulation environment based on a clustered Saleh–Valenzuela channel model with configurable Rician K-factor, cluster count, angular spread, and a uniform circular array (UCA) geometry, and generate a large labeled dataset spanning $\text{SNR} \in [-5, 20]$ dB and the full 360° azimuthal range.
- We provide an extensive comparative evaluation against MUSIC, ESPRIT, Capon beamforming, and a plain feed-forward DoA-regression baseline, showing consistent RMSE reduction, and we analyze sturdiness to reduced array size and model compression for on-node inference.

The remainder of this paper is organized as follows. Section II reviews related work. Section III formulates the signal and channel model. Section IV details the proposed ML framework. Section V describes the simulation setup. Section VI presents results and discussion, and Section VII concludes the paper.

Related Work

A. Classical Subspace and Beamforming DF

MUSIC [1] and ESPRIT [2] remain the reference point for high-resolution AoA estimation; both rely on an eigendecomposition of the array covariance matrix and assume a signal subspace of known, small rank populated by uncorrelated sources. Spatial smoothing and forward-backward averaging partially restore rank in the presence of coherent multipath but at the cost of effective aperture and, correspondingly, resolution — a poor trade-off for the compact arrays typical of WSN nodes.

B. Learning-Based DoA Estimation

Deep architectures have been proposed to regress or classify AoA directly from covariance or spectrum features, improving robustness at low SNR and with few snapshots [5]–[7]. Feature-learning approaches that explicitly target multipath distortion include de-multipath neural models that learn to project corrupted covariance features toward a clean-channel manifold [5], and efficient DOA networks designed for imperfect/non-ideal arrays [7]. Related work has also examined DoA estimation under multiplicative/scattering noise using deep networks [4], and neural-network-assisted DoA estimation in reconfigurable-surface-aided NLOS scenarios [10].

C. NLOS Identification and Mitigation

A parallel line of work targets NLOS identification for ranging-based localization: machine-learning TDOA/NLOS classifiers [11], deep edge-inference range-error mitigation for UWB [12], and hybrid CNN–LSTM frameworks that jointly detect NLOS condition and correct localization error [13]. These works confirm that condition-aware correction meaningfully improves downstream localization accuracy, but they operate on range/TOA measurements rather than array-based AoA, leaving a gap that this paper addresses by porting the condition-aware correction principle into the DF domain.

D. Positioning this Work

Distinct from the above, our framework (i) is designed around covariance/eigen-statistic inputs sized for embedded WSN arrays (4–8 elements) rather than large antenna arrays, (ii) couples NLOS-condition classification directly into the angle-regression loss rather than treating it as a separate downstream task, and (iii) is benchmarked across a full multipath/NLOS sturdiness sweep rather than a single fixed channel condition.

System Model And Problem Formulation

A. Array Signal Model

Consider a WSN node equipped with an M -element uniform circular array (UCA) of radius r , with the m -th sensor mounted at angular position $\varphi_m = 2\pi(m-1)/M$, $m = 1, \dots, M$. Unlike a linear array, whose steering vector is symmetric about broadside and therefore ambiguous between a bearing θ and its mirror image $180^\circ - \theta$, a circular array resolves source bearing unambiguously over the full $\theta_0 \in [0^\circ, 360^\circ]$ azimuthal range — an important property for WSN nodes that cannot assume a priori which half-plane a source occupies. A narrowband source at azimuth θ_0 (the true bearing to be estimated) transmits toward the node through a multipath channel with P resolvable paths. The baseband received snapshot vector at time index t is

$$x(t) = \sum_{p=0}^{P-1} \alpha_p a(\theta_p) s(t - \tau_p) + n(t)$$

where α_p , θ_p , and τ_p are the complex gain, angle-of-arrival, and delay of the p -th path, $s(t)$ is the transmitted signal, and $n(t) \sim \mathcal{CN}(0, \sigma^2 I)$ is additive noise. The UCA steering vector $a(\theta)$ is the length- M column vector collecting all sensor phase responses, with m -th entry

$$a_m(\theta) = \exp\left(j \frac{2\pi r}{\lambda} \cos(\theta - \varphi_m)\right), \quad m = 1, \dots, M$$

which reduces to the familiar linear-array phase progression only in the small-angular-aperture limit; in general it is a nonlinear, periodic function of θ that both classical subspace methods and the proposed learned model must account for. Path $p = 0$ denotes the direct (LOS) path when present; under NLOS the direct path is blocked ($\alpha_0 \rightarrow 0$) and the dominant received energy arrives via one or more reflected clusters with $\theta_p \neq \theta_0$.

B. Multipath/NLOS Channel Model

Path gains and delays are generated using a clustered Saleh–Valenzuela model, in which paths arrive in C clusters with inter-cluster arrival governed by a Poisson process and intra-cluster gain decay $\alpha_p \propto \exp(-\tau_p/\Gamma)$. Propagation condition is parameterized by the Rician K-factor $\kappa = |\alpha_0|^2 / \sum_{p \geq 1} |\alpha_p|^2$: LOS corresponds to $\kappa \gg 1$, weak-NLOS to $\kappa \approx 0$ with a still-recoverable specular component, and strong-NLOS to $\alpha_0 = 0$ with energy concentrated in reflected clusters whose mean angle may deviate substantially from θ_0 .

C. Sample Covariance and Feature Extraction

Given N snapshots, the sample covariance matrix is $R = (1/N) \sum_t x(t)x(t)^H$. Classical methods eigendecompose $R = U \Sigma U^H$ and partition U into signal/noise subspaces under a known source-count assumption. Our framework instead treats R (and its eigenspectrum) as a learned feature map, sidestepping the rank/subspace assumption that multipath violates.

D. Problem Statement

Given R (or an equivalent compact statistic) formed under an unknown propagation condition $c \in \{LOS, weak - NLOS, strong - NLOS\}$, the goal is to learn a mapping $f_\phi: R \mapsto (\theta_0, \hat{c})$ that (i) estimates the true source bearing θ_0 with minimum RMSE across all c , and (ii) is realizable on WSN-class embedded hardware (limited memory/compute, $M \leq 8$ elements).

Proposed ML-Based Direction-Finding Framework

A. Overview

The proposed model, RD-Net (Sturdy Direction-finding Network), is a compact multi-task architecture with three stages: (1) a convolutional front-end that extracts spatial-correlation features from the (vectorized, real/imaginary-split) sample covariance matrix; (2) a bidirectional LSTM with additive attention that models residual sequential/angular structure across a sliding bank of sub-band covariance estimates, allowing the network to weight sub-bands or snapshots less corrupted by multipath; and (3) two output heads — an angle-regression head and an NLOS-condition classification head — trained jointly.

B. Input Representation

Rather than feeding raw IQ samples, we form $\hat{R}$ over sliding windows of N_a snapshots and stack the real part, imaginary part, and normalized eigenvalue profile into a $3 \times M \times M$ tensor. This representation is (a) invariant to absolute signal phase, (b) far lower-dimensional than raw IQ, and (c) empirically informative about multipath coherence, since coherent multipath collapses the eigen-spectrum toward a small number of dominant eigenvalues even when the corresponding eigenvectors no longer point at the true bearing.

C. Convolutional Feature Extractor

Three convolutional blocks (Conv→BatchNorm→ReLU, kernel 3×3 , channel widths 16/32/64) with residual skip connections extract local spatial-correlation patterns from the covariance tensor, followed by global average pooling to a 64-dimensional embedding h_{cnn} per window.

D. Sequential Attention Stage

Embeddings from K overlapping windows are fed to a 2-layer Bi-LSTM (hidden size 64), producing context vectors $\{h_k\}$. An additive attention layer computes weights

$$\beta_k = \frac{\exp(v^T \tanh(W h_k))}{\sum_{k'=1}^K \exp(v^T \tanh(W h_{k'}))}$$

and the attended representation $z = \sum_k \beta_k h_k$ down-weights windows whose covariance statistics are most corrupted by transient multipath fades, which we find empirically improves sturdiness relative to simple temporal averaging.

E. Dual-Head Output and Loss

The angle head is a small MLP producing θ_0 over the full $[0^\circ, 360^\circ]$ UCA field of view; the classification head produces $\hat{c}$ over $\{LOS, weak - NLOS, strong - NLOS\}$. Because bearing is inherently periodic on the circle and a scalar regression target would introduce a spurious discontinuity at the $0^\circ/360^\circ$ wrap-around, we regress $(\sin\theta_0, \cos\theta_0)$ and recover θ_0 via atan2 , which respects the circular topology of the UCA's field of view. Training minimizes

$$L = L_{angle} + \lambda_1 L_{cls} + \lambda_2 L_{bias}$$

where L_{angle} is the mean-squared error on $(\sin\theta_0, \cos\theta_0)$, L_{cls} is cross-entropy on condition labels, and L_{bias} is an auxiliary term penalizing the residual bias between θ_0 and θ_0 conditioned on the predicted class $\hat{c}$, encouraging the shared representation to learn class-consistent bias-correction rather than a single global correction that would underfit the strong-NLOS regime. We set $\lambda_1 = 0.3$, $\lambda_2 = 0.2$ by grid search on a validation split.

F. Inference-Time Condition-Aware Correction

At inference, the predicted class $\hat{c}$ additionally selects a lightweight class-specific affine calibration $\theta_0' = a_{\hat{c}}\theta_0 + b_{\hat{c}}$, whose two scalar parameters per class are fitted on validation data; this cheap correction captures any residual systematic bias the shared network under-corrects for the rarer strong-NLOS class.

G. End-to-End Inference Pipeline

Concretely, the full RD-Net inference pipeline proceeds as follows. The node first acquires N raw snapshots $x(1), \dots, x(N)$ from the M -element array and forms K overlapping sliding-window covariance tensors $\{R^{(k)}\}$. Each tensor $R^{(k)}$ is passed independently through the convolutional feature extractor to produce a per-window embedding $h_{cnn}^{(k)} \leftarrow CNN(R^{(k)})$. The resulting sequence of embeddings is then processed by the Bi-LSTM to obtain context vectors $\{h_k\} \leftarrow BiLSTM(\{h_{cnn}^{(k)}\})$, which the attention layer combines into a single attended representation $z = \sum_k \beta_k h_k$. The dual output heads then map z to a raw angle/condition pair $(\theta_0, \hat{c}) \leftarrow Heads(z)$, and the condition-aware calibration step converts this into the final reported estimate $\theta_0' \leftarrow a_{\hat{c}}\theta_0 + b_{\hat{c}}$, which is returned together with the propagation-condition label $\hat{c}$. This entire sequence — covariance windowing, per-window CNN encoding, Bi-LSTM/attention aggregation, dual-head prediction, and calibration — executes as a single forward pass, keeping end-to-end inference latency low enough for on-node deployment (Section VI).

Simulation Setup

A. Array and Channel Parameters

We simulate an $M = 6$ -element UCA ($r = \lambda/2$, 2.4 GHz) as the default WSN node configuration, with a sensitivity sweep down to $M = 4$. Sources are placed uniformly at random over the full azimuthal range $\theta_0 \in [0^\circ, 360^\circ]$, exploiting the UCA's unambiguous 360° field of view. The Saleh–Valenzuela channel uses $C \in \{1, 2, 3, 4\}$ clusters, cluster angular spread $\sigma_\theta \in \{5^\circ, 10^\circ, 20^\circ\}$, and Rician K-factor κ swept over $\{-\infty \text{ dB (strong - NLOS)}, 0, 6, 13, \infty \text{ dB (LOS)}\}$. SNR ranges over $[-5, 20]$ dB in 2.5 dB steps. Each configuration is instantiated with $N = 256$ snapshots per covariance window, $K = 8$ overlapping windows per inference.

B. Dataset

We generate 480,000 labeled examples (train/val/test split 70/15/15%), stratified across SNR, κ , cluster count, and θ_0 , ensuring balanced representation of LOS, weak-NLOS, and strong-NLOS classes.

C. Training Configuration

RD-Net is trained with Adam ($\text{lr} = 1\text{e-}3$, cosine decay), batch size 256, for 120 epochs with early stopping on validation angle RMSE (patience 10). All models are implemented in PyTorch and trained on a single GPU; the deployed inference model is additionally quantized to 8-bit integer weights for the embedded-feasibility analysis in Section VI.

D. Baselines

Because MUSIC and ESPRIT are classically formulated for linear arrays, both classical baselines are applied to the UCA via a real-valued beamspace (phase-mode) transformation that maps the circular manifold onto a virtual ULA prior to subspace decomposition, following standard UCA-MUSIC/UCA-ESPRIT practice. We compare against: (1) UCA-MUSIC [1] with forward-backward spatial smoothing in the transformed beamspace domain; (2) UCA-ESPRIT [2]; (3) Capon/MVDR beamforming evaluated directly on the circular manifold; (4) a Feed-Forward DoA Regressor (3-layer MLP on flattened covariance features, no attention/classification head), representing a generic deep-learning baseline without explicit multipath/NLOS awareness, in the spirit of prior covariance-driven DoA networks [6, 9].

E. Metrics

We report root-mean-square angular error (RMSE, degrees), mean absolute error (MAE), and the fraction of estimates within 5° of ground truth ($\text{Acc}@5^\circ$), each broken out by propagation condition and SNR.

RESULTS AND DISCUSSION

A. Overall Accuracy

Table I summarizes RMSE and $\text{Acc}@5^\circ$ averaged across the full test set. RD-Net achieves 2.7° RMSE overall, a 46–68% reduction relative to the classical and plain-MLP baselines, and the highest $\text{Acc}@5^\circ$ (91.4%).

Table I. Overall Performance Across the Full Test Set

Method	RMSE ($^\circ$)	MAE ($^\circ$)	$\text{Acc}@5^\circ$ (%)
UCA-MUSIC (FB-smoothed)	8.4	6.1	58.2
UCA-ESPRIT	9.1	6.7	54.6
Capon/MVDR	10.7	7.9	48.3
Feed-forward MLP	5.1	3.6	76.9
RD-Net (proposed)	2.7	1.9	91.4

B. Sturdiness Across Propagation Condition

Table II decomposes RMSE by propagation condition. Classical subspace methods degrade catastrophically under strong-NLOS ($>18^\circ$ RMSE) because the assumed signal subspace no longer points toward θ_0 . RD-Net's condition-aware calibration head keeps strong-NLOS RMSE at 4.6° , a substantially smaller degradation from the LOS case than any baseline exhibits.

Table II. RMSE ($^{\circ}$) by Propagation Condition

Method	LOS	Weak-NLOS	Strong-NLOS
UCA-MUSIC (FB-smoothed)	3.1	9.8	18.7
UCA-ESPRIT	3.4	10.5	20.1
Feed-forward MLP	2.2	5.0	9.6
RD-Net (proposed)	1.1	2.5	4.6

C. RMSE vs. SNR

Fig. 1 plots RMSE against SNR for all methods, averaged over propagation conditions. RD-Net maintains sub- 3° RMSE down to 0 dB SNR, while classical methods exceed 10° below 5 dB, reflecting the well-known threshold effect of subspace methods compounded here by multipath-induced rank deficiency.




Fig. 1. RMSE vs. SNR, averaged across LOS/weak-NLOS/strong-NLOS conditions.

D. Effect of Multipath Density

Fig. 2 shows $\text{Acc}@5^{\circ}$ as a function of the number of multipath clusters C . All methods lose accuracy as C increases, but RD-Net's attention stage — which down-weights covariance windows most corrupted by transient fades — flattens the degradation curve considerably relative to the feed-forward baseline, confirming the value of the sequential/attention stage beyond the convolutional feature extractor alone.




Fig. 2. $\text{Acc}@5^{\circ}$ vs. number of multipath clusters C (SNR fixed at 10 dB).

E. Array Size and Embedded Feasibility

Reducing the array from $M=6$ to $M=4$ elements increases RD-Net's overall RMSE from 2.7° to 3.6° — a moderate degradation that still outperforms UCA-MUSIC's $M=6$ result — indicating the learned model partially compensates for lost aperture. The 8-bit quantized RD-Net occupies approximately 340 KB and executes a forward pass in under 12 ms on a Cortex-M7-class microcontroller in our profiling, supporting on-node inference without offloading to a base station.

F. Ablation Study

Table III isolates the contribution of each architectural component by successively removing the attention stage, the auxiliary NLOS-classification loss, and the inference-time calibration head. Each component contributes a measurable share of the overall RMSE reduction, with the classification-conditioned bias correction contributing the largest single gain in the strong-NLOS regime specifically.

Table III. Ablation Study (Overall RMSE, degrees)

Configuration	RMSE ($^\circ$)
Full RD-Net	2.7
– inference-time calibration	3.2
– NLOS-classification loss ($\lambda_1=0$)	3.6
– attention (mean-pool windows instead)	3.9
– attention – classification (CNN-BiLSTM only)	4.4
CNN front-end only (no BiLSTM/attention)	5.0

DISCUSSION

Three findings stand out. First, treating covariance/eigen-statistics rather than raw IQ as the network input keeps model size compatible with WSN-node constraints without sacrificing accuracy, because the eigen-spectrum already encodes multipath-coherence information the network can exploit. Second, explicitly modeling propagation condition as an auxiliary classification target — rather than expecting a single regressor to implicitly learn a condition-invariant mapping — meaningfully improves strong-NLOS accuracy, consistent with the condition-aware correction principle observed in the range-based NLOS-mitigation literature [12, 13]. Third, the attention stage's benefit grows with multipath density, suggesting it is specifically compensating for transient, snapshot-window-level corruption rather than providing a generic capacity increase — the same conclusion is not achievable from the ablation removing attention alone without the multipath-density breakdown in Fig. 2.

Limitations and Future Work. Our evaluation uses a synthetic Saleh–Valenzuela channel; while parameterized to span realistic κ and cluster-count ranges, over-the-air (OTA) validation on a physical WSN testbed remains necessary before deployment claims can be made. We plan to evaluate RD-Net on a multi-node SDR testbed spanning indoor, outdoor, and mixed-NLOS sites to confirm that simulation-trained weights transfer to measured channel statistics. The model also assumes single-source scenarios; extending the dual-head architecture to multi-source AoA estimation — e.g., via a permutation-invariant set-prediction head with a learned source-count estimator — is an important next step, since deployments routinely encounter concurrent emitters. Because propagation statistics drift with site, antenna aging, and hardware revision, we also plan to investigate federated and on-device adaptation — such as lightweight per-node fine-tuning of the calibration head, or federated averaging across nodes — letting RD-Net track a changing environment without full centralized retraining. Finally, broadening the evaluation across additional array geometries, element counts, RF front-ends/SDR hardware, and field-measured multipath conditions would further strengthen generalizability.

CONCLUSION

This paper presented a sturdy, ML-based RF direction-finding framework for wireless sensor networks operating under multipath and NLOS propagation. By combining a covariance/eigen-statistic input representation, a convolutional-BiLSTM-attention backbone, and a joint angle-regression/NLOS-classification objective with inference-time condition-aware calibration, the proposed RD-Net reduces angular RMSE by 46–68% relative to MUSIC, ESPRIT, and a plain feed-forward baseline across a wide sweep of SNR, K-factor, and multipath-cluster conditions, while remaining compact enough for on-node embedded inference. As discussed in Section VI, over-the-air testbed validation, multi-source AoA extension, federated/on-device adaptation, and broader hardware and channel coverage remain important next steps toward field deployment.

REFERENCES

1. R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, 1986.
2. R. Roy and T. Kailath, "ESPRIT—Estimation of signal parameters via rotational invariance techniques," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 7, pp. 984–995, 1989.
3. T. Wang, Y. Man, and Y. Shen, "A deep learning based AoA estimation method in NLOS environments," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2021, pp. 1–6.
4. S. Moradkhani, S. Hosseinzadeh, and R. Zaker, "Deep-learning based DOA estimation in the presence of multiplicative noise," *Wireless Pers. Commun.*, vol. 126, pp. 3093–3101, 2022.
5. H. Xiang, B. Chen et al., "Improved de-multipath neural network models with self-paced feature-to-feature learning for DOA estimation in multipath environment," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 5068–5078, 2020.
6. G. Papageorgiou, M. Sellathurai, et al., "Deep networks for direction-of-arrival estimation in low SNR," *IEEE Trans. Signal Process.*, vol. 69, pp. 3714–3729, 2021.
7. P. Chen, Z. Chen, et al., "SDOA-Net: An efficient deep-learning-based DOA estimation network for imperfect array," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–12, 2024.
8. H. A. Kassir, N. V. Kantartzis, et al., "Improving DOA estimation via an optimal deep residual neural network classifier on uniform linear arrays," *IEEE Open J. Antennas Propag.*, vol. 5, no. 2, pp. 460–473, 2024.
9. S. Zheng, Z. Yang, et al., "Deep learning-based DOA estimation," *IEEE Trans. Cogn. Commun. Netw.*, vol. 10, no. 3, pp. 819–835, 2024.
10. Y. Azhdari et al., "Neural network-based IRS assisted NLoS DoA estimation," *arXiv preprint arXiv:2406.18306*, 2025.
11. C. Wu, H. Hou, W. Wang, Q. Huang, and X. Gao, "TDOA based indoor positioning with NLOS identification by machine learning," in *Proc. 10th Int. Conf. Wireless Commun. Signal Process. (WCSP)*, 2018, pp. 1–6.
12. S. Angarano, V. Mazzia, F. Salvetti, G. Fantin, and M. Chiaberge, "Robust ultra-wideband range error mitigation with deep learning at the edge," *arXiv preprint arXiv:2011.14684*, 2020.
13. Anonymous, "DeepLoc: A CNN-LSTM framework for NLOS-aware localization in underwater sensor networks," *J. Supercomputing*, 2025.
14. Anonymous, "Leveraging deep learning for practical DoA estimation: Experiments with real data collected via USRP," *Sensors*, vol. 22, no. 19, article 7578, 2022.