Mission-Critical ISAC via Quantum-XAI: Adaptive Waveform Design and Counterfactual Stability

Raza Ahmed, Muneeb Ahmad, Soo Young Shin

razaahmed@kumoh.ac.kr, muneeb.ahmad@kumoh.ac.kr, wdragon@kumoh.ac.kr

Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea

Abstract

This study proposes a closed-loop Quantum Neural Network-Explainable AI (QNN-XAI) controller for adaptive chirp selection in Chirp-Indexed AFDM (CI-AFDM) MIMO-ISAC links over mission-critical doubly dispersive channels. A deterministic probing column produces compact delay-Doppler statistics that form a low-dimensional controller state. Conditioned on this state, the QNN outputs a calibrated categorical distribution over pre-/post-chirp pair indices, enabling shortlist-based reception with $M = 2^{b_1}! \ll !Q$ candidates instead of an exhaustive Q-hypothesis scan. Explainability is embedded in the loop: attribution-based masking suppresses noise-dominated features, counterfactual perturbations quantify boundary proximity (decision stability), and an uncertainty-triggered safeguard activates conservative fallback under adverse conditions. Simulations show that QNN-XAI improves BER, reduces decision latency (T_{dec}), increases P_D at fixed P_{FA}, and enhances joint reliability $R(\epsilon, \tau)$, while lowering per-frame chirp-selection overhead.

1. Introduction

Sixth-generation (6G) networks are expected to support mission-critical connectivity under stringent latency, reliability, and energy-efficiency requirements, often in dynamic settings where the controller observes only partial and noisy evidence [1]. This has accelerated the adoption of learning-enabled operation across the wireless stack, yet practical deployment still hinges on decisions that can be inspected, stress-tested, and trusted. In parallel, quantum computing paradigms motivate compact decision models that can express richer nonlinear mappings under constrained inference budgets, making quantum neural networks (QNNs) attractive for time-critical control loops [2].

Explainable AI (XAI) provides a route to accountable learning in wireless systems. However, much of the current approach remains open-loop: explanations are generated after an action is taken and rarely inform the next control decision [3]. This mismatch becomes acute for chirp-indexed AFDM-based ISAC over doubly dispersive channels, where exhaustive search across large chirp banks can dominate decision time and computational cost [4]. To address this, we develop a closed-loop QNN-XAI controller for CI-AFDM in which a deterministic probing column compresses each frame into low-dimensional delay-Doppler summaries, a QNN ranks pre-/post-chirp candidates to shrink residual testing from Q hypotheses to a shortlist $M \ll Q$, and an in-loop counterfactual boundary-proximity signal triggers conservative fallbacks when probe evidence is ambiguous. Using the same probe stream for controlled-P_{FA} sensing, the proposed loop improves both BER and P_{D} under a fixed false-alarm design in mission-critical doubly dispersive operation.

2. System Model

We consider a MIMO-ISAC transceiver with N_t transmit and N_r receive antennas operating over a doubly dispersive delay-Doppler channel. Each frame t contains K CI-AFDM columns of length N, where the first column is a known probing column used for sensing and for extracting compact delay-Doppler descriptors, and the remaining $K-1$ columns carry data. CI-AFDM applies a candidate pre-/post-chirp pair (indexed by $q \in \{1, \ldots, Q\}$) before and after the IDAFT, yielding a structured waveform family that adapts to channel dispersion. After prefix removal and standard preprocessing, the receiver forms a stacked observation $\mathbf{y}_t$ and, for each candidate index q, constructs an effective channel $\mathbf{H}_{\text{eff},t}(q)$ together with an MMSE estimate $\widehat{\mathbf{x}}_t(q)$. The exhaustive oracle decision selects the chirp-pair index by minimizing the residual error:

$$q_t^{\text{orc}} = \arg \min_{q \in \{1,\ldots,Q\}} \left\| \mathbf{y}_t - \mathbf{H}_{\text{eff},t}(q)\widehat{\mathbf{x}}_t(q) \right\|_2^2, \qquad (1)$$

which is accurate but expensive when Q is large. The proposed closed-loop controller uses the probing column to generate a low-dimensional state vector $\mathbf{z}_t$ and outputs a calibrated distribution over indices; only the most plausible candidates are evaluated, reducing per-frame decision time while retaining the original CI-AFDM index mapping.

This research jointly evaluates communication and sensing outcomes. Let $\widehat{\mathbf{H}}_t$ denote the channel estimate used by the receiver at the frame t (in the chosen processing domain). Channel estimation quality is reported via:

$$\mathrm{NMSE}_t \triangleq \frac{\mathbb{E}\left[\|\mathbf{H}_t - \widehat{\mathbf{H}}_t\|_F^2\right]}{\mathbb{E}\left[\|\mathbf{H}_t\|_F^2\right]}. \tag{2}$$

Although channel-estimation quality directly drives downstream equalization accuracy and thus influences BER, sensing reliability, and even the detection latency, this work does not report NMSE explicitly because its impact is already reflected in the end-to-end metrics shown in the following section. For sensing, the probing column yields a scalar score s_t and we decide $\widehat{\mathcal{H}}_t = \mathbb{I}\{s_t > \tau\}$ to report the operating point through:

$$P_{\mathrm{FA}}(\tau) = \Pr(s_t > \tau \mid \mathcal{H}_t = 0), \tag{3}$$

$$P_{\mathrm{D}}(\tau) = \Pr(s_t > \tau \mid \mathcal{H}_t = 1). \tag{4}$$

On the control side, the controller maps (optionally refined) probe descriptors to a categorical policy $\pi_\theta(q \mid \mathbf{z}_t)$ and selects a shortlist I_t of size $M \ll Q$ containing the top-M indices by probability:

$$I_t = \mathrm{Top}_M\left(\{\pi_\theta(q \mid \mathbf{z}_t)\}_{q=1}^{Q}\right), \tag{5}$$

so the residual test in (1) is computed only for $q \in I_t$. The captured posterior mass tracks the selection reliability:

$$\kappa_t \triangleq \sum_{q \in I_t} \pi_\theta(q \mid \mathbf{z}_t), \tag{6}$$

Where a high value κ_t indicates a low probability of discarding strong candidates, the decision latency is ultimately defined as the elapsed wall-clock time required to finalize the chirp-pair selection. Since residual evaluation predominantly influences this process, T_{dec} scales approximately with the number of evaluated indices, implying a near-proportional reduction when substituting Q evaluations with M.

3. Simulation Results

Figure 1 shows an apparent reduction in BER with increasing SNR across all schemes, with the learning-based controllers consistently outperforming the MMSE baseline. Conventional AI improves BER by reducing mismatched chirp-pair choices. At the same time, XAI provides an additional gain by suppressing noisy delay-Doppler cues and prioritizing the dominant factors that drive the decision. QNN-XAI achieves the lowest BER across the full SNR range, with the most visible separation in the medium-to-high SNR regime (around the 10^{-3} region), where its curve exhibits a distinct left shift relative to XAI and conventional AI, indicating more reliable candidate ranking and fewer residual evaluation errors.

Figure 2 highlights the computational advantage of QNN-XAI. Conventional AI exhibits the highest latency, while XAI gradually reduces it as SNR increases. In contrast, QNN-XAI shows a sharp latency drop after the mid-SNR region, consistent with a more concentrated policy that enables aggressive selection-based evaluation.

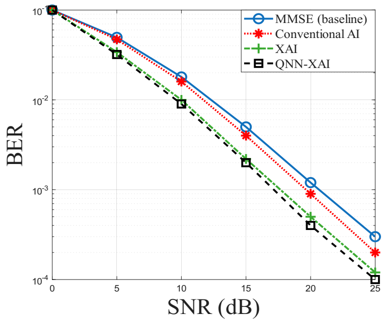

Figure 1. BER comparison of QNN-XAI vs conventional approaches.

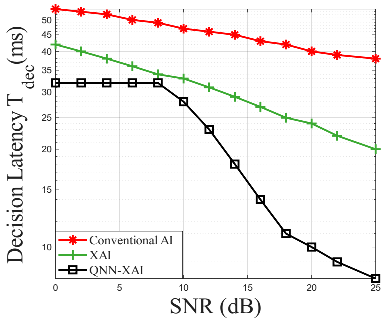

Figure 2. Decision latency of QNN-XAI, AI, and conventional XAI.

4. Conclusion

This paper proposes a closed-loop QNN-XAI controller for selecting chirps in CI-AFDM MIMO-ISAC over doubly dispersive channels. The probe-derived delay-Doppler summaries create a compact state, and the QNN gives a calibrated distribution to narrow down the list of candidates, which cuts down on the Q-hypothesis oracle search. In-loop explainability (attribution masking, counterfactual/uncertainty signals) makes decisions more stable and allows for a safe fallback. Simulations demonstrate improved BER and P_D at a constant P_{FA}, and increased reliability at lower SNR compared to MMSE, AI, and XAI.

5. Acknowledgment

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00553810, 50%) This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program(IITP-2025-RS-2024-00437190, 50%) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation)

References

[1] S. Mahboob, J. Dai, A. Soysal, and L. Liu, “Transforming future 6G networks via O-RAN-empowered NTNs,” *IEEE Comm. Mag*, 2025.

[2] R. Ahmed, S. M. H. Abidi, and S. Y. Shin, “Quantum Representation Learning for Images via Disentangling Quantum Autoencoders,” *한국통신학회 학술대회논문집*, pp. 821–822, 2025.

[3] S. Jere, L. Zheng, K. Said, and L. Liu, “Towards xAI: Configuring RNN weights using domain knowledge for MIMO receive processing,” *IEEE Trans. Wireless Commun.*, 2025.

[4] S. Ahmad, M. Ahmad, and S. Y. Shin, “AFDM-PPCIM: Pre-and-Post Chirp Index Modulation for Delay-Doppler Resilience in Terrestrial and Spaceborne Communications,” *한국통신학회 학술대회논문집*, pp. 244–245, 2025.