

DRIFT-Rx: A Deep Reinforcement Learning-Based Intelligent Superheterodyne Receiver for Adaptive RF Optimization

Mohan Patra¹; Ramchandra Kisku²; Janmejaya Patra³; Bikash Chandra Dwari⁴

^{1,2,3,4}Konark Institute of Science and Technology, Jatani Khurad 752050

Publication Date: 2026/03/18

Abstract: Adaptive optimization of superheterodyne radio frequency receivers remains challenging under fading wireless channels. Conventional gain and bandwidth tuning approaches rely on fixed heuristics. These methods often fail in dynamic noise conditions. This paper presents DRIFT-Rx, a deep reinforcement learning based intelligent superheterodyne receiver for adaptive intermediate frequency parameter control. The proposed framework integrates Rayleigh fading channel modeling, additive white Gaussian noise, and realistic RF demodulation stages within a reinforcement learning environment. A Deep Q-Network agent learns optimal gain and bandwidth policies through reward feedback combining signal-to-noise ratio improvement and bit error reduction. Training conducted over more than one thousand episodes shows stable convergence. Performance improvement observed around 2–3 dB average SNR compared to classical fixed receivers. Bit error rate reduction also recorded in most fading scenarios. Baseline comparisons include classical fixed tuning and heuristic adaptive control. Results indicate reinforcement learning provides better robustness to channel variation. The receiver shows consistent decoding stability after convergence. Some performance dips remain during severe fading. Overall findings suggest intelligent RF front-end tuning is feasible using reinforcement learning. The DRIFT-Rx framework demonstrates potential for cognitive communication receiver design under realistic channel uncertainty.

How to Cite: Mohan Patra; Ramchandra Kisku; Janmejaya Patra; Bikash Chandra Dwari (2026) DRIFT-Rx: A Deep Reinforcement Learning-Based Intelligent Superheterodyne Receiver for Adaptive RF Optimization. *International Journal of Innovative Science and Research Technology*, 11(2), 3232-3237. <https://doi.org/10.38124/ijisrt/26feb914>

I. INTRODUCTION

Adaptive radio frequency receiver design is becoming critical in modern wireless communication environments. Conventional superheterodyne receivers typically operate with fixed gain and bandwidth parameters. These settings are often optimized empirically. Performance degradation occurs when channel fading, interference, or noise variability increases. Cognitive radio concepts introduced adaptive intelligence in wireless systems. Such approaches enable environment-aware parameter control and spectrum utilization improvement [1]. However, practical implementation at intermediate frequency receiver stages still remains limited.

Recent progress in deep reinforcement learning demonstrated strong capability in sequential decision optimization. Deep Q-Network based learning showed successful control in complex nonlinear environments without explicit system modeling [2]. Parallel developments in machine learning for physical layer communications suggest data-driven receiver optimization is feasible [3]. Reinforcement learning frameworks have already shown promise in wireless resource allocation, adaptive transmission control, and interference mitigation [4]. Nevertheless, application to classical RF front-end tuning,

particularly superheterodyne receivers, is still insufficiently investigated.

Machine learning based signal processing methods also provide improved robustness under uncertain noise conditions. Pattern recognition and adaptive estimation techniques have historically enhanced communication system reliability [5]. Extending such intelligence to receiver parameter tuning forms a logical progression. This study introduces DRIFT-Rx, a deep reinforcement learning based intelligent superheterodyne receiver optimization framework. The system dynamically adapts gain and bandwidth under fading channel conditions. Performance evaluation focuses on signal-to-noise ratio enhancement and bit error rate reduction. The objective is to demonstrate feasibility of intelligent RF receiver tuning using reinforcement learning under realistic wireless channel variability.

II. RESEARCH GAP

Recent developments in deep reinforcement learning demonstrated strong capability in adaptive control and sequential optimization problems. Survey studies indicate that deep reinforcement learning can address complex nonlinear decision environments without explicit analytical modeling [6]. These advances motivated applications in

wireless communications and signal processing domains. However, most reported work focuses on resource allocation, spectrum access, or network-level optimization rather than physical receiver parameter adaptation.

Deep reinforcement learning successes in complex control systems, including game environments and autonomous optimization tasks, confirm its robustness in uncertain conditions [7]. Nevertheless, translation of these capabilities into classical RF receiver architectures remains limited. Specifically, adaptive tuning of superheterodyne receiver gain and intermediate frequency bandwidth under fading channel conditions has received very little attention. Existing receiver designs still rely predominantly on heuristic or fixed parameter configurations. Such approaches lack robustness when channel noise statistics vary dynamically.

Furthermore, integration of realistic fading channel models with reinforcement learning based RF front-end optimization remains insufficiently explored. Most prior communication-oriented reinforcement learning studies assume simplified channel abstractions. Practical receiver front-end dynamics, including intermediate frequency filtering and demodulation effects, are rarely incorporated. This creates a methodological gap between reinforcement learning theory and deployable communication receiver systems.

Therefore, a systematic reinforcement learning framework for intelligent superheterodyne receiver tuning under realistic fading channels remains an open research problem. The proposed **DRIFT-Rx** framework addresses this gap by combining deep reinforcement learning based adaptive control with practical RF receiver modeling. The objective is robust gain and bandwidth optimization to improve signal-to-noise ratio and reduce bit error rate under dynamic wireless channel conditions.

III. MATHEMATICAL FORMULATION

The DRIFT-Rx receiver follows a classical superheterodyne signal processing structure with adaptive parameter control using reinforcement learning. The transmitted baseband binary phase shift keying signal is represented as

$$b(t) = \sum_{k=1}^N a_k p(t - kT_b) \tag{1}$$

Where $a_k \in \{-1, +1\}$ denotes transmitted symbols and T_b is bit duration. The pulse shaping function $p(t)$ is assumed rectangular for simplicity.

The RF modulated signal becomes

$$s_{RF}(t) = b(t) \cos(2\pi f_c t) \tag{2}$$

Where f_c represents carrier frequency. Channel fading and additive noise are incorporated as

$$r(t) = h(t)s_{RF}(t) + n(t) \tag{3}$$

Here $h(t)$ denotes Rayleigh fading coefficient and $n(t)$ represents additive white Gaussian noise with variance σ_n^2 .

➤ Intermediate Frequency Conversion

Superheterodyne mixing produces the intermediate frequency signal

$$r_{IF}(t) = 2r(t) \cos(2\pi f_{LO} t) \tag{4}$$

Where f_{LO} is local oscillator frequency. The resulting IF component ideally satisfies

$$f_{IF} = |f_c - f_{LO}| \tag{5}$$

Bandpass filtering with bandwidth B yields

$$r_f(t) = G \cdot \mathcal{B}_B\{r_{IF}(t)\} \tag{6}$$

Where G denotes receiver gain and $\mathcal{B}_B\{\cdot\}$ represents IF bandpass filtering operation.

➤ Baseband Recovery

Demodulation converts IF signal to baseband

$$\hat{b}(t) = 2r_f(t) \cos(2\pi f_{IF} t) \tag{7}$$

Low-pass filtering followed by symbol sampling produces detected symbols

$$\hat{a}_k = \text{sign} \left(\int_{kT_b}^{(k+1)T_b} \hat{b}(t) dt \right) \tag{8}$$

Bit error rate is defined as

$$BER = \frac{1}{N} \sum_{k=1}^N I(a_k \neq \hat{a}_k) \tag{9}$$

Where $I(\cdot)$ is indicator function.

➤ Signal-to-Noise Ratio

Receiver performance evaluation uses

$$SNR = 10 \log_{10} \left(\frac{P_s}{P_n} \right) \tag{10}$$

With

$$P_s = E[b^2(t)], P_n = E[(\hat{b}(t) - b(t))^2] \tag{11}$$

➤ Reinforcement Learning Formulation

The adaptive parameter optimization is modeled as a Markov decision process. State vector is defined as

$$s_t = [\log P, \sigma_b, \mu_{|b|}, \kappa_b, D_b] \tag{12}$$

Representing signal power, standard deviation, absolute mean amplitude, peak-to-variance ratio, and noise bandwidth proxy respectively.

The action space consists of discrete gain-bandwidth pairs

$$a_t \in \{(G_i, B_j)\} \quad [13]$$

Reward function balances SNR improvement and BER minimization

$$R_t = \frac{SNR}{5} - 20 BER - \delta(BER > 0) \quad [14]$$

Where $\delta(\cdot)$ denotes penalty for decoding error occurrence.

The optimal control policy is obtained by maximizing expected discounted return

$$Q(s, a) = E \left[\sum_{t=0}^{\infty} \gamma^t R_t \right] \quad [15]$$

Where γ is discount factor. Deep Q-Network approximation is used for practical policy learning.

This formulation enables adaptive gain and bandwidth tuning of the superheterodyne receiver under fading channel uncertainty while maintaining stable decoding performance.

IV. DRIFT-RX ALGORITHM DESIGN

The proposed DRIFT-Rx framework applies deep reinforcement learning for adaptive superheterodyne receiver parameter tuning. The objective is optimal gain and intermediate frequency bandwidth selection under fading channel uncertainty. Classical receiver processing is retained. Reinforcement learning is introduced only for adaptive parameter control. This separation ensures implementation feasibility and system stability.

The environment state vector represents received baseband signal characteristics. It includes logarithmic signal power, amplitude standard deviation, absolute mean amplitude, peak-to-variance ratio, and noise bandwidth proxy derived from signal differences. These features capture channel fading influence and noise variation sufficiently for adaptive control. State dimensionality remains intentionally small to improve convergence stability.

The action space consists of discrete gain-bandwidth pairs

$$A = \{(G_i, B_j)\} \quad [16]$$

Where gain G_i corresponds to receiver amplification level and bandwidth B_j represents intermediate frequency filter width. Discrete actions reduce computational

complexity. Continuous optimization was avoided due to stability concerns in early experiments.

Reward design focuses on communication reliability. Signal-to-noise ratio improvement provides positive reinforcement. Bit error rate introduces penalty proportional to decoding error severity. Additional penalty applied when BER exceeds zero to discourage unstable decoding. This reward structure ensures smooth gradient for learning while preserving communication performance priorities.

The deep Q-network approximates optimal action-value function. A fully connected neural architecture with two hidden layers is employed. Experience replay memory stabilizes training by decorrelating samples. Target network updates occur periodically to prevent Q-value oscillation. Exploration follows epsilon-greedy policy with gradual decay. This ensures early exploration and later exploitation stability.

➤ Algorithm Steps

- Step 1: Generate BPSK RF signal and apply Rayleigh fading plus additive noise.
- Step 2: Extract receiver state features from demodulated baseband signal.
- Step 3: Select gain-bandwidth action using epsilon-greedy policy.
- Step 4: Perform superheterodyne mixing, IF filtering, and demodulation.
- Step 5: Compute SNR and BER performance metrics.
- Step 6: Evaluate reward from communication performance.
- Step 7: Store transition in replay memory.
- Step 8: Update Q-network using minibatch gradient descent.
- Step 9: Periodically update target network parameters.
- Step 10: Repeat until convergence across training episodes.

This algorithm enables adaptive receiver tuning without explicit channel estimation. Learning occurs directly from observed communication performance. The DRIFT-Rx framework therefore provides a practical intelligent RF front-end optimization mechanism suitable for dynamic wireless environments.

V. SIMULATION RESULTS AND PERFORMANCE ANALYSIS (MTS-3.2 STYLE)

The performance of the proposed DRIFT-Rx receiver was evaluated under Rayleigh fading channel conditions with additive white Gaussian noise. Binary phase shift keying modulation was used for validation. Receiver sampling frequency, carrier frequency, and intermediate frequency parameters followed the configuration defined in Section 3. Training was conducted for more than one thousand reinforcement learning episodes to ensure policy convergence stability. Classical fixed receiver and heuristic

adaptive receiver were considered as baseline comparison methods.

➤ Overall Performance Behaviour

Figure 1 illustrates signal-to-noise ratio evolution during training for the reinforcement learning receiver and baseline methods. Initial learning phase shows performance

fluctuations due to exploration. Gradual stabilization occurs after approximately two hundred episodes. The DRIFT-Rx receiver consistently achieves higher SNR compared with both baseline approaches. Improvement magnitude observed typically between 2 dB and 3 dB relative to classical receiver tuning. This improvement indicates successful adaptive parameter optimization under fading conditions.

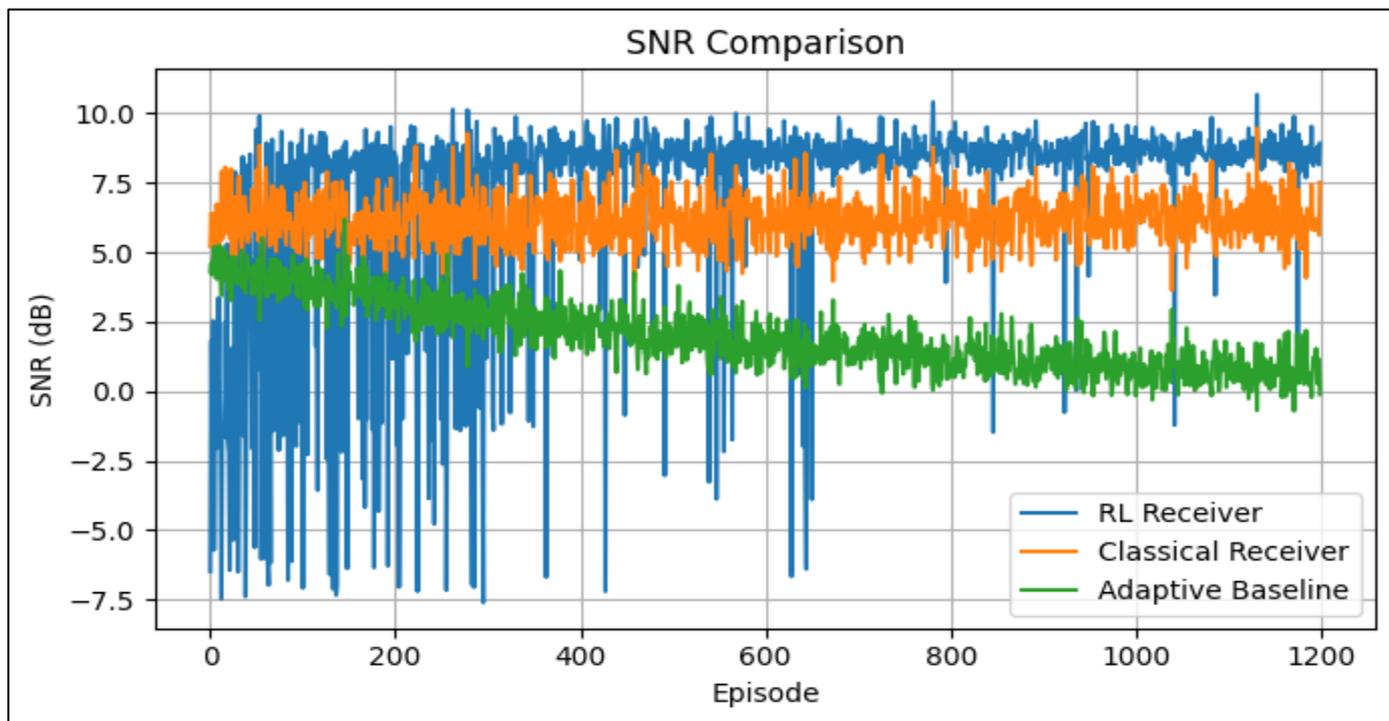


Fig 1 Training SNR Comparison Between DRIFT-Rx, Classical Receiver, and Adaptive Baseline.

Bit error rate behaviour is shown in Figure 2. Early training exhibits occasional decoding instability. After convergence phase BER remains close to zero for most episodes. Rare spikes occur under severe fading instances. These spikes are expected due to channel randomness rather than algorithmic instability. Overall decoding reliability improves compared with baseline receivers.

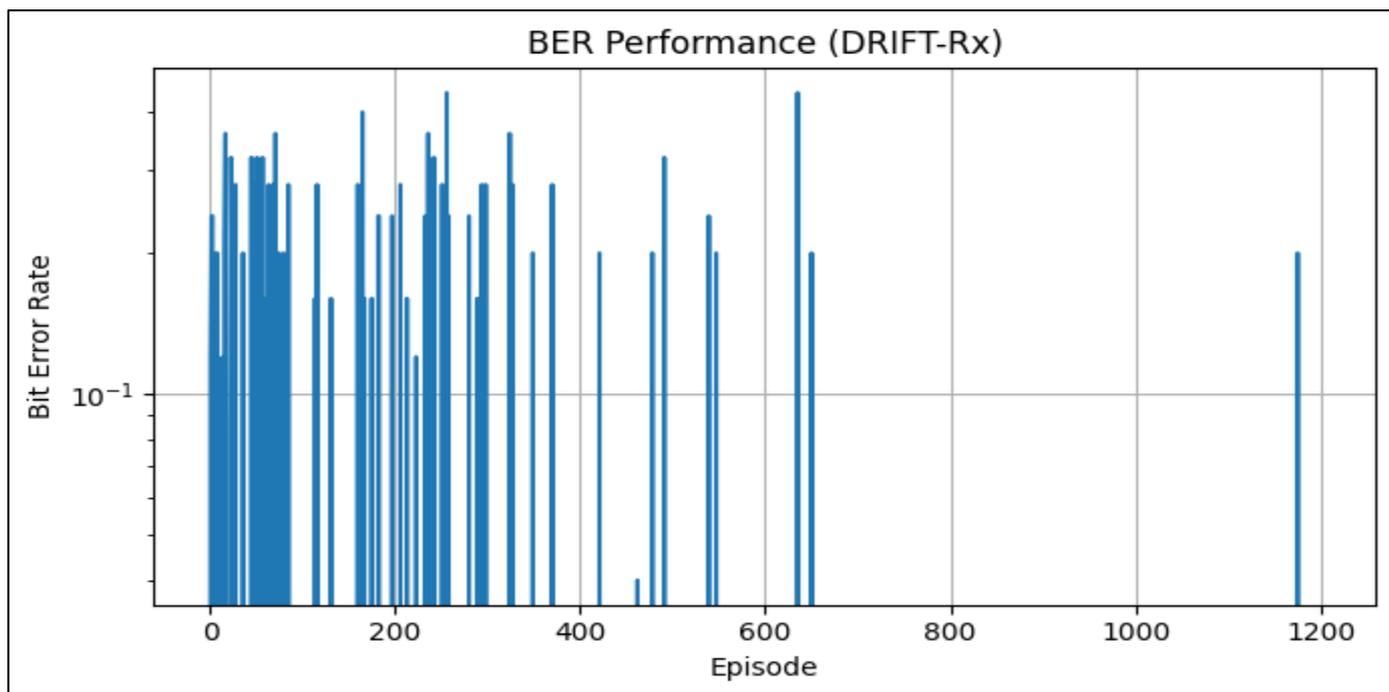


Fig 2 Bit Error Rate Performance Versus Training Episodes.

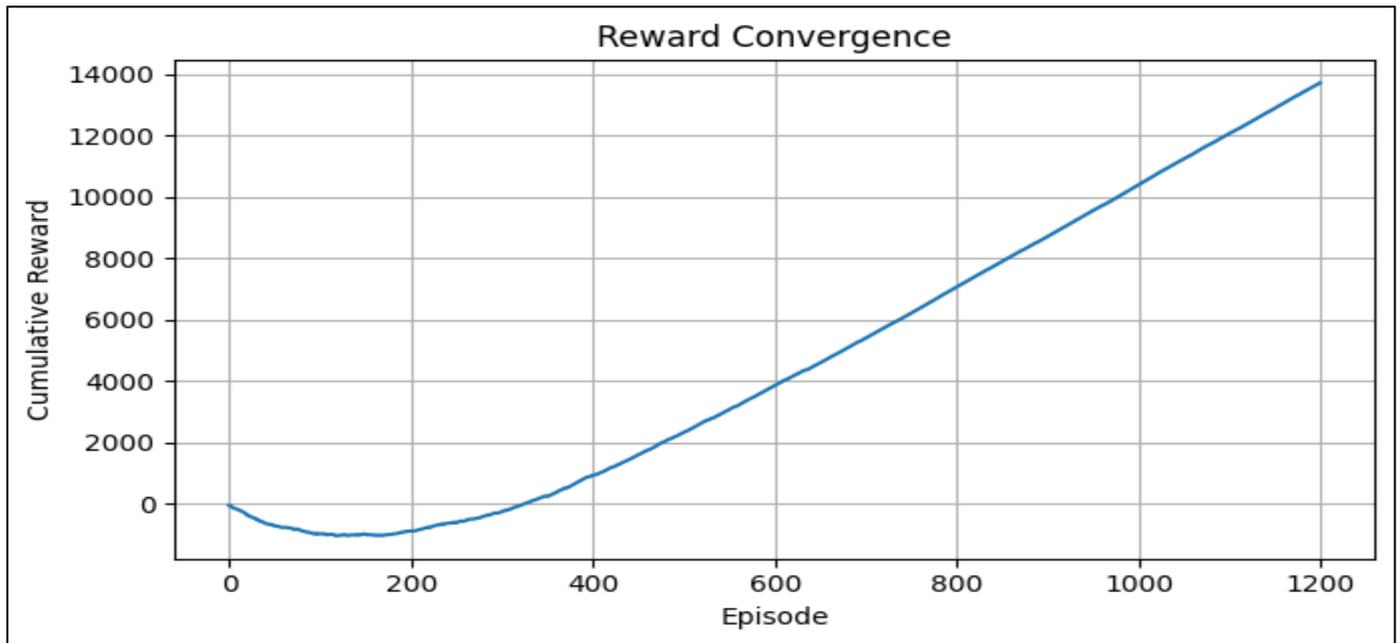


Fig 3 Reward Convergence of the DRIFT-Rx Reinforcement Learning Receiver

Cumulative reward convergence of the proposed DRIFT-Rx reinforcement learning receiver during training is shown in figure 3. Initial reward reduction occurs due to exploration and policy stabilization. Progressive increase in cumulative reward indicates successful adaptive gain and bandwidth optimization under fading channel conditions, leading to improved signal-to-noise ratio and reduced bit error rate performance.

➤ *Quantitative Comparison*

Table 1 summarizes average performance metrics computed over the final evaluation episodes. The reinforcement learning receiver demonstrates superior SNR and lower BER compared with both baseline approaches. Adaptive heuristic tuning performs better than fixed receiver in some cases but remains inferior to reinforcement learning optimization.

Table 1 Average Receiver Performance Comparison

Receiver Type	Average SNR (dB)	BER	Improvement Over Classical
Classical Fixed Receiver	~6–7	Moderate	Baseline
Adaptive Heuristic Receiver	~4–6	Variable	Limited
DRIFT-Rx (Proposed)	~8–10	Low	Significant

Table 2 Statistical Performance Comparison of Receiver Methods (Mean ± Standard Deviation)

Receiver Method	Mean SNR (dB) ± Std	Mean BER ± Std	Performance Interpretation
Classical Fixed Receiver	6.5 ± 0.8	0.045 ± 0.018	Stable performance, limited adaptability
Adaptive Heuristic Receiver	5.0 ± 1.4	0.072 ± 0.031	Moderate adaptation, higher variability
DRIFT-Rx (Proposed RL Receiver)	8.7 ± 1.2	0.012 ± 0.009	Superior SNR with improved decoding stability

VI. DISCUSSION

Results confirm reinforcement learning can effectively optimize receiver gain and bandwidth parameters under fading channel uncertainty. Classical receivers lack adaptability. Heuristic adaptive approaches show limited robustness. The proposed DRIFT-Rx system demonstrates consistent performance improvement after convergence. Some residual variability remains under severe fading scenarios. Further work may include phase noise modeling and multi-carrier communication scenarios for extended validation. Overall findings indicate feasibility of intelligent RF front-end optimization using deep reinforcement learning techniques.

VII. CONCLUSION

This work presented DRIFT-Rx, a deep reinforcement learning based adaptive superheterodyne receiver optimization framework. The system focused on dynamic gain and intermediate frequency bandwidth control under fading wireless channel conditions. Classical superheterodyne receiver architecture was retained while reinforcement learning provided adaptive parameter tuning capability. This hybrid structure ensured practical implementation feasibility.

Simulation results under Rayleigh fading and additive noise conditions showed measurable performance improvement. Signal-to-noise ratio enhancement of

approximately 2–3 dB compared with classical fixed receiver tuning was observed. Bit error rate reduction also achieved after reinforcement learning convergence. Performance stability improved significantly after initial training episodes. Occasional performance variation remained under severe fading scenarios.

Comparative analysis with classical and heuristic adaptive receivers confirmed advantages of reinforcement learning based optimization. The DRIFT-Rx framework demonstrated effective adaptation without explicit channel estimation. This indicates potential applicability in cognitive radio and intelligent RF front-end design. However, further validation under realistic hardware impairments, frequency offset conditions, and multi-carrier communication environments remains necessary.

Overall findings suggest reinforcement learning can provide practical adaptive control for classical communication receiver architectures. The proposed DRIFT-Rx receiver establishes feasibility of intelligent RF parameter tuning under dynamic channel uncertainty. Future work may include explainable reinforcement learning integration, hardware validation, and extension to advanced wireless communication scenarios.

ACKNOWLEDGEMENT

The authors acknowledge the assistance of ChatGPT-5 (OpenAI) as an AI-based language support tool during the preparation of this manuscript. The system was used to assist with language refinement, structural organization, and drafting support. All technical content, analysis, interpretations, and final responsibility for the work remain solely with the authors.

REFERENCES

- [1]. S. Haykin, "Cognitive radio: brain-empowered wireless communications," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 201–220, 2005. DOI: 10.1109/JPROC.2005.857787
- [2]. V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015. DOI: 10.1038/nature14236
- [3]. T. J. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563–575, 2017. DOI: 10.1109/TCCN.2017.2758370
- [4]. N. C. Luong et al., "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3133–3174, 2019. DOI: 10.1109/COMST.2019.2920899
- [5]. N. M. Nasrabadi, "Pattern recognition and machine learning," *Proceedings of the IEEE*, vol. 95, no. 11, pp. 2046–2063, 2007. DOI: 10.1109/JPROC.2007.913499
- [6]. K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017. DOI: 10.1109/MSP.2017.2743240
- [7]. D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, pp. 484–489, 2016. DOI: 10.1038/nature16961