Robust Backscatter Detection via Deep Reinforcement Learning Against an Evolving Adversarial Jammer

Hao Hoang Tran, Quang Tuan Do, Tung Son Do, Dongwook Won, Junsuk Oh, and Sungrae Cho
School of Computer Science and Engineering
Chung-Ang University
Seoul 06974, South Korea
Email: {hhtran, dqtuan, tsdo, dwwon, jsoh}@uclab.re.kr, srcho@cau.ac.kr

***Abstract*—Bistatic backscatter systems are a key enabler for energy-efficient IoT, but their covert operation is threatened by the rise of intelligent adversaries. While prior work has focused on designing covert transmission strategies against simple, static detectors, the resilience of an intelligent detector against an adaptive adversary remains underexplored. This paper addresses this gap from the perspective of an AI-powered warden. We model the warden as a Deep Reinforcement Learning (DRL) agent trained to dynamically adjust its detection threshold. To create a formidable threat, we introduce an evolving adversarial jammer, modeled as a generator network, which is trained using a black-box optimization technique—Evolution Strategies (ES)—to produce maximally deceptive noise signals. We evaluate the DRL-based warden's performance across several scenarios, including no jamming, Gaussian noise jamming, and our proposed evolving ES-based jammer. Simulation results demonstrate that while the evolving ES-jammer poses a significant threat that actively suppresses the warden's performance, the DRL agent successfully learns a robust counter-policy. It proves its resilience by maintaining a stable detection accuracy significantly above the random guess baseline, highlighting the necessity of adaptive, learning-based defenses in modern adversarial communication environments.**

***Index Terms*—Covert communication, bistatic backscatter, deep reinforcement learning (DRL), evolution strategies (ES), adversarial jamming, physical layer security.**

I. Introduction

Bistatic backscatter communication has emerged as a key enabler for energy-efficient Internet of Things (IoT) networks, allowing passive tags to transmit data by reflecting carrier signals [1]. However, the open nature of the wireless medium exposes these transmissions to adversaries, or "wardens," necessitating robust security measures. Beyond traditional encryption, covert communication, which aims to hide the very existence of a transmission, provides a fundamental layer of security [2].

Much of the existing research in this domain focuses on designing sophisticated transmission or jamming strategies for the legitimate user to remain undetected by a simple, static warden [1], [3]. A key limitation of such work is the assumption of a non-adaptive adversary. This assumption creates a significant vulnerability, as it fails to account for a modern, intelligent adversary that can learn and adapt its strategy over time. This raises a critical question: how can a detector maintain its effectiveness when the adversary is also intelligent and actively evolving its attack strategy?

This paper addresses this challenge by shifting the focus from the transmitter to the receiver, framing the problem from the perspective of an AI-powered warden that must ensure robust detection. We model an adversarial arms race between two learning agents, a paradigm that reflects the growing application of advanced DRL algorithms like Deep Q-Networks (DQN) and Deep Deterministic Policy Gradient (DDPG) in complex IoT network environments [5], [7]. The main contributions of this work are as follows:

- We model the Warden as an intelligent Deep Reinforcement Learning (DRL) agent. Instead of using a fixed detection rule, our Warden learns a dynamic policy to adjust its detection threshold in real-time based on the observed signal characteristics.
- We design a formidable, adaptive adversary: an Evolving Adversarial Jammer. This jammer is modeled as a generator network trained with a black-box optimization technique known as Evolution Strategies (ES) to continuously refine its jamming waveform to be maximally deceptive.
- We demonstrate through simulation that our DRL-based Warden learns a robust and adaptive policy. While its performance is challenged by the evolving ES-based jammer, it successfully maintains a stable and positive accuracy margin, highlighting the necessity of adaptive defenses against intelligent adversaries.

II. System Model and Adversarial Problem Formulation

A. System and Channel Model

We consider a bistatic backscatter system operating in a hostile environment, as depicted in Fig. 1. The system comprises a Carrier Emitter (CE), a passive Tag, a legitimate Reader, and our intelligent Warden. We introduce an external, adaptive adversary, the Evolving Adversarial Jammer, whose objective is to transmit a confusing artificial noise (AN) signal, $z(i)$, to prevent the Warden from detecting the Tag's transmissions.

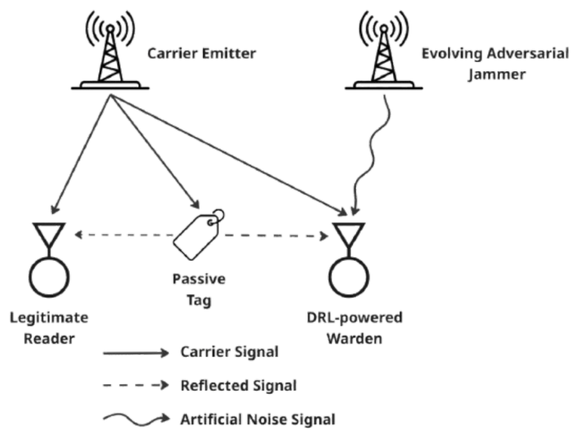

Fig. 1. System model for robust detection. The DRL-powered Warden must learn to detect the Tag's faint reflected signal in the presence of a sophisticated, Evolving Adversarial Jammer.

The Warden faces a binary hypothesis testing problem to determine if the Tag is active (H_1) or silent (H_0). The signal received at the Warden, $y_w(i)$, is given by:

$$y_w(i) = \begin{cases} h_{cw}c(i) + h_{jw}z(i) + n_w(i), & \text{under } H_0 \\ \sqrt{\beta}h_{ct}f_{tw}s(i)c(i) + h_{cw}c(i) & \\ \quad + h_{jw}z(i) + n_w(i), & \text{under } H_1 \end{cases} \tag{1}$$

where $c(i)$ is the carrier signal from the CE, $s(i)$ is the Tag's symbol, β is the reflection coefficient, and $n_w(i)$ is thermal noise. The term $z(i)$ represents the sophisticated jamming signal from the adaptive adversary.

To create a challenging and realistic simulation environment, we implement a specific channel model. The direct path from the CE to the Warden is modeled as a stable, line-of-sight channel with a fixed gain ($h_{cw} = 1.0$). In contrast, all other paths, including the backscatter path (h_{ct}, f_{tw}) and the jamming path (h_{jw}), experience independent Rayleigh fading, with their gains drawn from a complex Gaussian distribution for each signal sample. Furthermore, to encourage the Warden to learn a truly robust policy, the environment operates in a "stationary realism" mode. In each step of the simulation, the power of both the backscatter signal and the jamming signal are randomly scaled, forcing the Warden to adapt to a constantly changing Signal-to-Noise and Interference Ratio.

B. The Adversarial Arms Race

We formulate the interaction between the two intelligent agents as a continuous, zero-sum "arms race," where each agent's objective is to counter the other. The protagonist of our framework is the DRL-Warden, modeled as a Deep Reinforcement Learning (DRL) agent. Its goal is to learn a detection policy, π_W, that maps its observation of the environment (the state, s_t) to an optimal action, a_t (the choice of a detection threshold, τ_t). The agent's performance is measured by its ability to maximize its long-term cumulative reward, which is based on the accuracy of its H0/H1 classification. The Warden's objective is to learn a policy that is robust enough to maintain high accuracy even as the jammer's strategy evolves.

Conversely, the antagonist is the ES-Jammer, modeled as a generator network G_J trained with Evolution Strategies (ES). Its goal is to learn an optimal generator, G_J^*, that produces a jamming waveform $z(i)$ which is maximally deceptive. In this arms race, the jammer's objective at any given time is to find a strategy that minimizes the Warden's current detection accuracy. This creates a challenging, non-stationary environment where the Warden cannot rely on static assumptions about the interference and must instead learn an adaptive and generalizable detection policy.

C. The DRL-Powered Warden

To create a resilient detector, we model the Warden as an intelligent agent trained with Deep Reinforcement Learning (DRL). DRL has been shown to be a powerful tool for optimizing complex surveillance and monitoring tasks in dynamic wireless environments [4], [6]. The core of our Warden's intelligence lies in its ability to process raw signal data directly. This follows a broader trend in wireless communications where deep learning models, particularly Convolutional Neural Networks (CNNs), are increasingly used as the foundation for complex tasks, ranging from semantic understanding [8] to the adversarial detection explored in this work. This allows the Warden to learn a dynamic decision-making policy to adjust its detection threshold, τ. To handle the high-dimensional, raw signal data, we design a hybrid architecture that combines a Convolutional Neural Network (CNN) for perception and a Deep Q-Network (DQN) for decision-making.

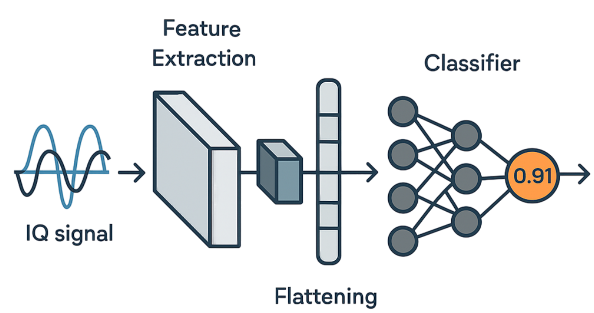

Fig. 2. Conceptual architecture of the Hybrid DRL-Warden. A CNN processes the raw signal into a feature vector, which serves as the state for a DQN that selects the optimal detection threshold.

The Warden's learning problem is formulated as a Markov Decision Process (MDP):

- **State ($\mathcal{S}$):** The state, s_t, is not the raw signal itself, but a compact feature vector extracted from a block of I/Q samples by a dedicated 1D CNN, as depicted in Fig. 2. This allows the agent to base its decisions on high-level patterns rather than noisy, low-level data.
- **Action ($\mathcal{A}$):** The action, a_t, is the selection of a detection threshold, τ_t, from a discrete set of ten possible power

levels. The Warden compares the received power of an incoming signal block to this threshold to decide between H_0 and H_1.
- **Reward ($\mathcal{R}$):** The agent receives a simple and direct reward for its actions: $r_t = +1.0$ for a correct detection and $r_t = -1.0$ for an incorrect detection. This reward structure directly incentivizes the agent to maximize its classification accuracy.

The Warden employs a DQN to learn its policy. The DQN is trained using experiences stored in a replay buffer to approximate the optimal action-value function, $Q^*(s, a)$. This enables it to learn a robust policy, $\pi_W(\tau|s)$, that maps the observed signal features to the optimal detection threshold, thereby maximizing its long-term accuracy.

D. The Evolving Adversarial Jammer

To provide a truly adaptive and challenging adversary, we model the jammer as a generative network, G_θ, parameterized by weights θ. This network outputs a complex noise waveform $z(i) = G_\theta(v)$ from a random latent vector v.

Critically, this generator is not trained with a conventional GAN setup. Instead, we employ Evolution Strategies (ES), a powerful, gradient-free optimization technique well-suited for non-differentiable or black-box objectives, such as fooling a separate DRL agent. The ES training algorithm, which underpins the jammer's ability to evolve, is outlined below:

1) A "population" of candidate generators is created by adding random Gaussian noise to the current generator's weights: $\theta_k = \theta + \sigma\epsilon_k$, where σ is the noise standard deviation.
2) The fitness of each perturbed generator, $F(\theta_k)$, is evaluated based on its ability to deceive a discriminator. This discriminator's detection rate on H1 signals, denoted as D_{rate}, serves as a proxy for the DRL-Warden's performance. The jammer's fitness is defined using a shaped reward function to provide a stronger learning signal:
$$F(\theta_k) = -\log(D_{\text{rate}}(G_{\theta_k})), \tag{2}$$
where G_{θ_k} is the generator network with the perturbed weights θ_k. A lower detection rate for the discriminator results in an exponentially higher fitness score for the jammer.
3) The central weights, θ, are updated by moving them in the direction of the "evolutionary gradient," which is calculated as a weighted sum of the perturbations, where each perturbation is weighted by its corresponding fitness score.

This process allows the jammer to continuously "evolve" its noise-generation strategy to find and exploit weaknesses in the Warden's current policy, creating a non-stationary and perpetually challenging environment.

E. Adversarial Training Methodology

The DRL-Warden and the ES-Jammer are trained in an interleaved loop, simulating a continuous arms race, as detailed in Algorithm 1. In this co-evolutionary process, the jammer periodically refines its strategy based on the warden's current performance, and the warden, in turn, continuously trains against this newly evolved.

1) The Jammer's Evolution: At a fixed interval, the ES-based Jammer performs its learning step. This phase represents the leader adapting its strategy. As a black-box optimization technique, ES "probes" the Warden's current policy by creating a population of slightly different noise generators. The fitness of each variant is determined by how successfully it deceives the Warden, measured empirically. By rewarding the perturbations that lead to a higher detection error for the Warden, the jammer effectively estimates an "evolutionary gradient" and updates its generator network to produce more deceptive signals.

2) The Warden's Adaptation: In every episode, the DRL-Warden performs its training step against the current, and potentially newly evolved, adversary. The Warden faces a "moving target"—a non-stationary environment where the statistics of the jamming noise can change. The Warden collects experiences into a large replay buffer, which is critical for stabilizing learning in this adversarial setting. By training on a diverse mix of both recent and past experiences, the Warden is prevented from overfitting to the jammer's latest strategy. The Warden's DQN is then updated by sampling from this buffer and minimizing the Bellman error via gradient descent. This step forces the Warden to learn a policy that is not just optimal for one type of noise, but is robust and generalizable enough to perform well against an evolving adversary.

III. Performance Evaluation

In this section, we present the numerical results to validate the robustness of our proposed DRL-based Warden. We evaluate its performance by training it within an adversarial arms race against an Evolving Adversarial Jammer and compare its learning dynamics to scenarios with simpler, static interference.

A. Simulation Setup

The experiment models a co-evolutionary "arms race" where the DRL-Warden and an ES-Jammer are trained in an interleaved loop.

1) The DRL-Warden: The Warden is a hybrid DRL agent designed for robust, learning-based detection.

- *Architecture:* Its "perception" module is a 1D Convolutional Neural Network ('WardenFeatureExtractor') that processes raw signal waveforms of length 128 into a feature vector. This CNN consists of two convolutional blocks, each using **ReLU** activation and max pooling. The "decision" module is a Deep Q-Network ('DQN'), a three-layer Multi-Layer Perceptron (MLP) with **ReLU** activations, which takes the feature vector as its state.
- *Action Space:* The DQN's output layer has 10 neurons, corresponding to 10 discrete detection thresholds spaced linearly from 0.1 to 2.5.

Algorithm 1 Adversarial Training of DRL-Warden and ES-Jammer

Initialize:
Warden's DQN parameters ψ.
Jammer's generator parameters θ.
Replay buffer $\mathcal{B}$ with capacity B_{max}.
Proxy discriminator D_{proxy} for fitness evaluation.
Learning rates α_W, α_J; population size N; update frequency K_{ES}.
for each training episode $e = 1, \dots, E$ **do**
  if $e \bmod K_{ES} == 0$ **then** {Evolve the Jammer}
  for $k = 1, \dots, N$ **do**
    Sample perturbation $\epsilon_k \sim \mathcal{N}(0, I)$.
    Create candidate weights $\theta_k = \theta + \sigma\epsilon_k$.
    Evaluate detection rate of the proxy discriminator:
    $D_{\text{rate}}(G_{\theta_k}) = \mathbb{E}_{v\sim P(v)}[D_{proxy}(G_{\theta_k}(v))]$.
    Calculate the fitness for the candidate:

$$F(\theta_k) = -\log(D_{\text{rate}}(G_{\theta_k})).$$

  end for
  Estimate the evolutionary gradient:

$$\nabla_\theta F \approx \frac{1}{N\sigma}\sum_{k=1}^{N} F(\theta_k)\epsilon_k.$$

  Update the central jammer parameters: $\theta \leftarrow \theta + \alpha_J \nabla_\theta F$.
  end if
  {Train the Warden for one episode}
  Observe initial state s_1 from the environment (using G_θ).
  for each step $t = 1, \dots, T$ **do**
    With probability ϵ, select random action a_t.
    Otherwise, select $a_t = \arg\max_a Q_\psi(s_t, a)$.
    Execute a_t, receive reward r_t and next state s_{t+1}.
    Store experience tuple (s_t, a_t, r_t, s_{t+1}) in $\mathcal{B}$.
    if $|\mathcal{B}| > B_{min}$ **then**
      Sample a mini-batch of M experiences $\{(s_j, a_j, r_j, s_{j+1})\}_{j=1}^{M}$ from $\mathcal{B}$.
      For each experience j, set the target Q-value:

$$y_j = r_j + \gamma \max_{a'} Q_\psi(s_{j+1}, a').$$

      Define the Bellman error loss function:

$$\mathcal{L}(\psi) = \frac{1}{M}\sum_{j=1}^{M}(y_j - Q_\psi(s_j, a_j))^2.$$

      Update the Warden's DQN parameters via gradient descent: $\psi \leftarrow \psi - \alpha_W \nabla_\psi \mathcal{L}(\psi)$.
    end if
  end for
end for

- *Training:* The DQN is trained using an ϵ-greedy policy, with epsilon decaying from 1.0 to 0.01 over the training episodes. Experiences are stored in a replay buffer of size 10,000. The network weights are updated by sampling batches of size 64 and minimizing the Mean Squared Error (Bellman error) loss using the **Adam optimizer** with a learning rate of `1e-4`.

2) The Evolving Adversarial Jammer: The jammer is a generative network (‘Generator‘) trained via Evolution Strategies (ES) to be a formidable, adaptive adversary.

- *Architecture:* The generator is a four-layer MLP with **LeakyReLU** activations and **LayerNorm** for stabilization. It takes a latent vector of dimension 64 as input and outputs a complex noise waveform of length 128.
- *Training (ES):* The jammer evolves every 5 episodes (‘ES_UPDATE_EVERY‘). In each evolution step, a population of 16 candidate generators (‘ES_POPSIZE‘) is created by perturbing the central weights with Gaussian noise. The fitness of each candidate is evaluated against a proxy detector (‘WardenDiscriminator‘) and be calculated using a shaped reward function (2). The central weights are then updated using the Adam optimizer along the estimated evolutionary gradient.

3) Environment and Scenarios: To ensure the Warden learns a generalizable and robust policy, the environment operates in a ”stationary realism” mode (‘REALISM_MODE = ”stationary”‘). For each signal sample generated, the power of both the artificial noise and the backscattered tag signal are randomly scaled by uniform multipliers drawn from the ranges [0.6, 1.3] and [0.6, 1.4], respectively. This forces the Warden to adapt to constantly changing signal conditions rather than overfitting to fixed power levels. We evaluate the Warden's performance across three distinct adversarial scenarios:

1) **No AN**: A baseline with only thermal noise.
2) **Gaussian AN**: A traditional adversary using statistically simple Gaussian noise.
3) **Generative AN (Ours)**: The most advanced adversary, our proposed ES-trained generator, which actively adapts its jamming signal every 5 episodes.

The key hyperparameters for the simulation are summarized in Table I.

B. Numerical Results and Discussion

The results of our adversarial training simulation are summarized in the four-panel dashboard in Fig. 3. This figure provides a comprehensive view of the co-evolutionary learning dynamics, tracking the performance of both the DRL-Warden and the ES-Jammer over 1,000 training episodes.

1) Warden's Detection Accuracy (Main Result): The top-left panel is the primary result of our experiment, showing the DRL-Warden's smoothed detection accuracy. The results clearly illustrate the Warden's adaptive learning capability and its resilience. In the ”No AN” baseline (orange curve), the Warden effectively learns the task, with its accuracy steadily rising and converging to the highest level of approximately

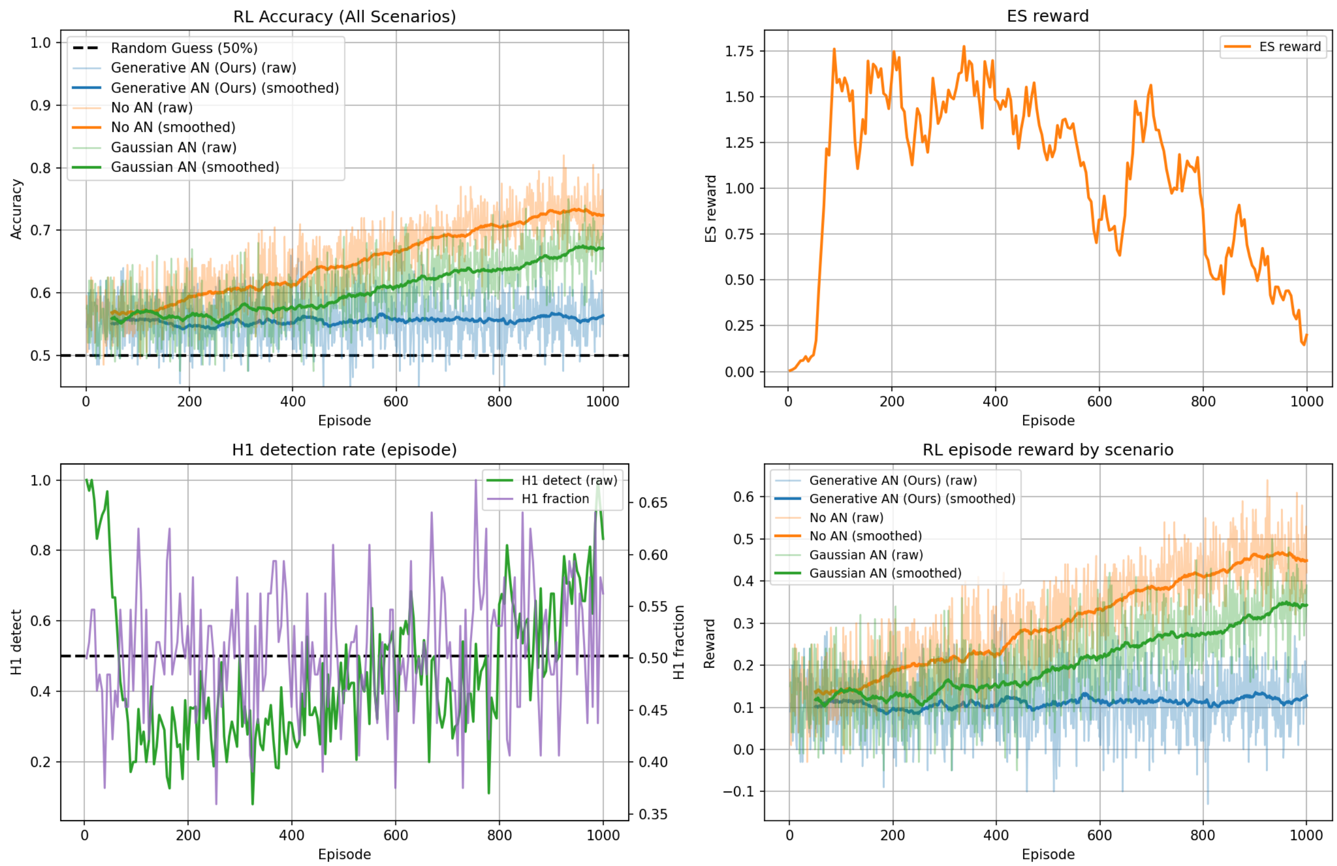

Fig. 3. Dashboard of adversarial training dynamics. (Top-Left) The DRL-Warden's smoothed detection accuracy across the three scenarios. (Top-Right) The fitness reward of the Evolving Adversarial Jammer over time. (Bottom-Left) Diagnostic plot of the jammer's H1 detection rate. (Bottom-Right) The raw and smoothed episode rewards for the DRL-Warden.

TABLE I
KEY SIMULATION PARAMETERS

Parameter	Value
DRL-Warden Parameters	
Training Episodes (`RL_EPISODES`)	1,000
Learning Rate (`RL_LR`)	1e-4
Replay Buffer Size	10,000
Action Space (Thresholds)	10 discrete levels
ES-Jammer Parameters	
Population Size (`ES_POPSIZE`)	16
Learning Rate (`ES_LR`)	0.05
Update Frequency (`ES_UPDATE_EVERY`)	5 Episodes
System Parameters	
Signal Length	128 samples
Backscatter Signal Power (`POWER_TAG`)	0.8
Batch Size	64

73%. When faced with a static "Gaussian AN" jammer (green curve), the task becomes more difficult, and the Warden's accuracy converges to a lower but still competent **67%**.

The most critical result is the "Generative AN (Ours)" scenario (blue curve), where the Warden faces the evolving ES-Jammer. Here, the Warden's learning is significantly suppressed. Its accuracy struggles to rise, hovering around **55-56%**, just above the 50% random guess baseline. This demonstrates that the ES-Jammer is a far more formidable adversary than static noise, as it actively adapts its strategy to counter the Warden's learning process. The fact that the Warden's accuracy does not collapse to 50% but instead maintains a small but consistent positive performance margin highlights its resilience.

2) The Adversarial Arms Race Dynamics: The top-right and bottom-right panels provide deeper insight into the adversarial arms race. The bottom-right panel shows the raw and smoothed rewards received by the DRL-Warden. The hierarchy of the smoothed lines directly mirrors the accuracy plot, confirming that the Warden receives the highest rewards in the easiest scenario ("No AN") and the lowest rewards when facing the most difficult adversary ("Generative AN").

Conversely, the top-right panel shows the fitness reward of the ES-Jammer over time. A higher reward for the jammer means it is being more effective at deceiving the Warden. The plot shows the jammer's reward spiking initially as it quickly finds exploits against the naive, untrained Warden. However, as the DRL-Warden begins to learn and adapt (around episode 600 onwards), we see a distinct downward trend in the jammer's reward. This is a crucial finding: it is direct evidence that the DRL-Warden is successfully learning a counter-strategy, making the jammer's task progressively harder and reducing its effectiveness. This demonstrates the Warden's robustness and its ability to fight back in the adversarial arms race.

3) Jammer's Deception Diagnostics: The bottom-left panel provides a diagnostic view of the jammer's performance during its fitness evaluation. The plot tracks the raw H1 detection rate by a proxy detector. The jammer's goal is to minimize this metric. The noisy but discernible trend corroborates the "ES reward" plot, showing the jammer's initial success followed by a struggle as the DRL-Warden learns and adapts.

In conclusion, the collective results from the dashboard paint a clear picture. While the evolving ES-Jammer poses a profound threat that significantly suppresses the Warden's performance, the DRL-Warden demonstrates remarkable resilience. It successfully learns a robust policy that prevents the jammer from completely dominating the engagement, maintaining a positive accuracy margin and actively pushing back against the adversary's adaptations. This highlights the critical necessity of employing adaptive, learning-based defenses like our DRL-Warden in modern, intelligent adversarial environments.

IV. Conclusion

In this paper, we shifted the paradigm of covert communications security from designing elusive transmitters to architecting a resilient, intelligent detector capable of operating against adaptive adversaries. We proposed a novel framework where the warden is modeled as a Deep Reinforcement Learning (DRL) agent that learns a dynamic detection policy. To rigorously test its capabilities, we introduced a formidable, evolving adversary: a generative jammer trained with Evolution Strategies (ES) to continuously refine its deceptive noise signals. Our simulations demonstrated that while the evolving adversary poses a significant threat that successfully suppresses detection performance, the DRL-based Warden proves its resilience by learning a robust policy. It successfully adapts to the jammer's changing strategy and maintains a stable detection accuracy, proving its ability to operate effectively in a non-stationary, hostile environment.

V. Acknowledgement

This work was supported by the IITP (Institute of Information & Communications Technology Planning & Evaluation) - ITRC (Information Technology Research Center) (IITP-2026-RS-2022-00156353, 50% / IITP-2026-RS-2023-00258639, 50%) grants funded by the Korea government (Ministry of Science and ICT).

References

[1] Y. Wang, S. Yan, W. Yang, Y. Huang, and C. Liu, "Energy-Efficient Covert Communications for Bistatic Backscatter Systems," *IEEE Trans. Veh. Technol.*, vol. 70, no. 3, pp. 2906-2910, Mar. 2021.

[2] B. A. Bash, D. Goeckel, and D. Towsley, "Limits of reliable communication with low probability of detection on AWGN channels," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 9, pp. 1921-1930, Sep. 2013.

[3] T. V. Sobers, B. A. Bash, S. Guha, D. Towsley, and D. Goeckel, "Covert communication in the presence of an uninformed jammer," *IEEE Trans. Wireless Commun.*, vol. 16, no. 9, pp. 6193-6206, Sep. 2017.

[4] W. J. Yun et al., "Cooperative Multiagent Deep Reinforcement Learning for Reliable Surveillance via Autonomous Multi-UAV Control," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 10, pp. 7086-7096, Oct. 2022.

[5] D. Kwon, J. Jeon, S. Park, J. Kim and S. Cho, "Multiagent DDPG-Based Deep Learning for Smart Ocean Federated Learning IoT Networks," *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 9895-9903, Oct. 2020.

[6] T. T. H. Pham, W. Noh, and S. Cho, "Multi-agent Reinforcement Learning based Optimal Energy Sensing Threshold Control in Distributed Cognitive Radio Networks with Directional Antenna," *ICT Express*, vol. 10, no. 3, pp. 472-478, June 2024.

[7] M. C. Ho, A.-T. Tran, D. Lee, J. Paek, W. Noh, and S. Cho, "A DDPG-based Energy Efficient Federated Learning Algorithm with SWIPT and MC-NOMA," *ICT Express*, vol. 10, no. 3, pp. 600-607, June 2024.

[8] T. S. Do, T. P. Truong, Q. T. Do, and S. Cho, "TranGDeepSC: Leveraging ViT Knowledge in CNN-based Semantic Communication System," *ICT Express*, vol. 11, no. 2, pp. 335-340, April 2025.