

Secure mmWave Beamforming with Proactive-ISAC Defense Against Beam-Stealing Attacks

Seyed Bagher Hashemi Natanzi ^{✉*}, Hossein Mohammadi ^{✉†}, Bo Tang ^{✉*}, Vuk Marojevic ^{✉†}

^{*}Electrical and Computer Engineering Department, Worcester Polytechnic Institute, USA

[†]Electrical and Computer Engineering Department, Mississippi State University, USA

Email: {snatanzi, btang1}@wpi.edu, {hm1125, vm602}@msstate.edu

Abstract—Millimeter-wave (mmWave) communication systems face increasing susceptibility to advanced beam-stealing attacks, posing a significant physical layer security threat. This paper introduces a novel framework employing an advanced Deep Reinforcement Learning (DRL) agent for proactive and adaptive defense against these sophisticated attacks. A key innovation is leveraging Integrated Sensing and Communications (ISAC) capabilities for active, intelligent threat assessment. The DRL agent, built on a Proximal Policy Optimization (PPO) algorithm, dynamically controls ISAC probing actions to investigate suspicious activities. We introduce an intensive curriculum learning strategy that guarantees the agent experiences successful detection during training to overcome the complex exploration challenges inherent to such a security-critical task. Consequently, the agent learns a robust and adaptive policy that intelligently balances security and communication performance. Numerical results demonstrate that our framework achieves a mean attacker detection rate of 92.8% while maintaining an average user SINR of over 13 dB.

Index Terms—6G, MIMO, Beamforming, Beam-Stealing, ISAC.

I. INTRODUCTION

mmWave communications provide high data rates for applications like AR/VR and connected vehicles [1]. This is achieved via highly directional beamforming, which mitigates severe path loss but introduces physical layer vulnerabilities [2]. Beam-stealing attacks, where adversaries hijack or eavesdrop beams, threaten link integrity and confidentiality [3]. Concurrently, ISAC is emerging as a key 6G feature, efficiently using shared resources for dual functionality [4]. ISAC's integration however also introduces new security threats.

Securing mmWave links against sophisticated beam-stealing attackers that may employ intelligent, adaptive strategies and exploit protocol knowledge to evade conventional defenses is a significant challenge. For instance, simple Power Delay Profile (PDP) analysis, although useful, can be circumvented by attackers capable of subtle manipulations [3]. This underscores the urgent need for robust, proactive, and adaptive defense mechanisms capable of countering dynamic and intelligent threats to ensure a reliable and trustworthy user experience.

Fig. 1 illustrates the threat and defense model considered in this work. A malicious attacker attempts to hijack the communication beam from a legitimate user. To counter this, we propose a novel framework where a DRL agent empowers

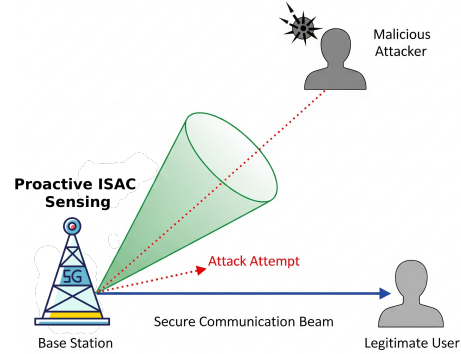


Fig. 1. System and defense overview. The base station maintains a secure communication link with the legitimate user while facing a beam-stealing attempt from a malicious attacker. The integrated ISAC module actively probes the environment, detects potential threats, and informs the DRL agent, which dynamically adjusts the beam direction and sensing effort accordingly.

the base station with proactive and adaptive beam control strategies. A key innovation is the use of ISAC not merely as a communication aid, but as an active sensing tool: the DRL agent leverages ISAC outputs to dynamically probe, detect, and respond to potential attacks. However, due to the complexity of the environment, standard agents often fail to discover secure policies. To address this, we introduce an intensive curriculum learning approach that ensures early successful detection experiences, guiding the agent toward robust convergence.

The primary contributions are:

- A security framework built on an advanced PPO agent for stable learning in a complex state-action space.
- A novel, intensive curriculum learning strategy that guarantees the agent experiences successful threat detection, a critical step to overcome the deep exploration challenges inherent in security-critical applications.
- A comprehensive analysis showing that the agent acquires a sophisticated, adaptive strategy that balances the detection rate with user communication quality.

This holistic approach significantly improves the resilience and security of mmWave links, ultimately enhancing the dependability of consumer oriented mmWave applications.

The remainder of this paper is organized as follows. Section II reviews the related work. Section III introduces the

system model, problem formulation, and overall framework. Section IV presents the proposed DRL agent and its integration with ISAC. Section V provides the simulation setup and performance evaluation. Finally, Section VII concludes the paper.

II. RELATED WORK

Beam-stealing attacks pose critical security threats in mmWave systems. Steinmetzer et al. [3] practically demonstrate such attacks on IEEE 802.11ad networks: By injecting forged feedback into the sector sweep protocol, they enabled Man-in-the-Middle (MITM) relays for eavesdropping. Addressing beam-stealing attacks without reliance on training data, Yang et al. [5] propose an image processing methodology. Their approach uses a Received Signal Strength Indicator (RSSI) map for joint detection and localization of multiple attackers, achieving 91% detection rates and sub-meter accuracy. Li et al. [2] propose SecBeam to counter sophisticated beam-stealing, specifically amplify-and-relay (AnR) attacks that can circumvent beacon authentication. This protocol analyzes the PDP to detect manipulated signal paths by verifying that legitimate, shorter paths exhibit stronger and earlier-arriving signals compared to potentially amplified, longer relay paths. Recent active attacks such as BeamCraft [6], which successfully manipulate Wi-Fi traffic by injecting forged Beamforming Feedback Information (BFI), demonstrate the critical need for proactive defenses against such clear-text feedback vulnerabilities.

Further exploring physical layer security, Qiu et al. [7] address hybrid threats in mmWave environments by introducing an artificial noise (AN)-aided robust multi-beam secure communication scheme. Their work centers on the joint design of information and AN beamforming to counteract coexisting active jammers and passive eavesdroppers, considering imperfect adversary Channel State Information (CSI), where the legitimate receivers employ Minimum Variance Distortionless Response (MVDR) for jamming suppression. In the context of ISAC, Xu et al. [8] combine physical layer covert transmission with ISAC functionalities. Their work designs transmit beamforming for both fully digital and hybrid architectures, enabling confidential, undetected communication to a covert user while concurrently supporting regular communication and sensing tasks, even under imperfect warden CSI. Furthermore, the security of Artificial Intelligence (AI) models for mmWave beamforming prediction is of critical concern. Kuzlu et al. [9] highlight the susceptibility of deep learning based beam predictors to adversarial attacks that corrupt input data and explore adversarial training and defensive distillation as mitigation techniques. While prior works focus on securing mmWave systems against attacks, optimizing performance in dynamic 6G environments is equally critical. Our recent work [10] uses DRL for adaptive beam switching to improve throughput and SNR in dynamic 6G environments. Similarly, Mohammadi et al. [11] employ multipath communications to counter 5G jamming attacks, enhancing physical layer security against physical layer threats.

Our DRL framework uniquely employs ISAC for proactive defense against mmWave beam-stealing. It distinctively optimizes secure beamforming and ISAC probing for enhanced situational awareness and robust and adaptive countermeasures.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. System Model

We consider a downlink mmWave ISAC scenario where a base station (BS) communicates with a legitimate user equipment (UE) while attempting to detect a potential beam-stealing attacker. The BS employs narrow directional beamforming to maximize link quality and simultaneously performs active sensing to identify suspicious targets in its environment.

Let $\mathbf{h} \in \mathbb{C}^N$ denote the channel vector between the BS and the UE, and let $\mathbf{w} \in \mathbb{C}^N$ represent the BS beamforming vector. The received signal is modeled as:

$$y = \mathbf{h}^H \mathbf{w} x + n,$$

where x is the transmitted symbol and n is complex Gaussian noise with variance σ^2 .

The SINR observed by the UE is:

$$\text{SINR} = \frac{|\mathbf{h}^H \mathbf{w}|^2}{\sigma^2}.$$

The BS dynamically adjusts its beam direction and ISAC sensing effort. The sensing module models the detection probability of the attacker as a function of sensing effort e_t and distance d_t :

$$P_d = 1 - e^{-\alpha e_t} e^{-\beta d_t},$$

where $\alpha, \beta > 0$ model sensing efficiency and distance attenuation.

Sensor measurements are subject to Gaussian errors in range and azimuth, and the agent observes a state vector \mathbf{s}_t that includes communication and sensing features.

B. Problem Formulation

The primary objective is to devise a control policy, $\pi(a_t | \mathbf{s}_t)$, that prioritizes security by maximizing attacker detection while maintaining adequate communication services. This is formulated as a DRL problem where the agent learns to maximize the expected cumulative discounted reward:

$$G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k}.$$

The specific structure of R_t is designed in the next section to align with the trade-off between detection, communication quality, and sensing effort.

IV. PROPOSED METHOD

Our proposed solution to optimize beamforming and sensing is an advanced DRL agent built upon the PPO algorithm. PPO is a state-of-the-art, policy-based method known for its stability and robust performance in complex, high-dimensional

environments. The agent is trained to dynamically adjust the base station's beam azimuth and ISAC effort to maximize security while maintaining communication quality. To overcome the significant exploration challenges inherent in this task, we introduce an intensive curriculum learning phase.

A. PPO Agent Design

The agent employs an Actor-Critic architecture, which consists of two separate neural networks:

- **The Actor Network:** This network learns the policy $\pi(a_t|s_t)$ by taking the current state as the input and outputting a probability distribution over the discrete action space.
- **The Critic Network:** This network estimates the value function $V(s_t)$, which predicts the expected cumulative reward from a given state. This value is used to assess the quality of the Actor's actions.

Both networks are implemented as fully connected multilayer perceptrons (MLPs). The input layer corresponds to the 7-dimensional state vector. This is followed by two hidden layers with 256 and 128 neurons, respectively, using ReLU activation functions. The Actor network has a final output layer with 5 neurons and a softmax activation to represent the action probabilities, while the Critic network has a single linear output neurons for the state value. The models are optimized using the Adam optimizer. PPO is chosen for its stability and efficiency in continuous and noisy environments. Unlike DQN, which requires discrete actions, PPO handles continuous control directly. Compared to A2C, PPO's clipping improves training stability, which is crucial in sparse-reward tasks like ours. Imitation Learning is unsuitable here due to the lack of expert demonstrations.

B. Reward Design

The reward function, R_t , at each step t , is designed as a weighted sum of key behavioral indicators to balance proactive defense and communication quality:

$$\begin{aligned} R_t = & w_{\text{det}} \cdot \mathbb{I}(\text{conf}_t > 0.7) \\ & + w_{\text{pro}} \cdot \mathbb{I}(\text{range}_t^{\text{true}} < 80\text{m} \wedge \text{effort}_t > 0.8) \\ & - w_{\text{unaware}} \cdot \mathbb{I}(\text{range}_t^{\text{true}} < 80\text{m} \wedge \text{conf}_t < 0.7) \\ & + w_{\text{com}} \cdot \mathbb{I}(\text{SINR}_t > 5\text{dB}), \end{aligned} \quad (1)$$

where $\mathbb{I}(\cdot)$ is an indicator function. The weights prioritize security: $w_{\text{det}} = 150$, $w_{\text{pro}} = 25$, $w_{\text{unaware}} = 5$, and $w_{\text{com}} = 0.5$. This structure encourages confident detections, rewards proactive sensing, penalizes ignorance of close threats, and preserves minimum SINR for service quality.

C. Training with Intensive Curriculum Learning

A primary challenge in this problem is the vast and sparse reward landscape where the agent may not find the high-reward states corresponding to a successful detection. Initial experiments without a curriculum confirm this, resulting in a 0% detection rate as the agent settles in a suboptimal policy of only maximizing the SINR. We design an intensive, two-phase curriculum learning strategy to solve this issue.

1) Phase 1: Forced Success Curriculum (First 1500 Episodes)

The goal of this initial phase is to guarantee that the agent experiences successful detection, thereby learning the value of security-oriented actions. In each of the first 1500 training episodes, a "forced success" mechanism selects five unique, random time steps and overrides the agent's actions to set the beam azimuth to the attacker's true direction and ISAC effort to its maximum value of 1.0, ensuring 100% guaranteed detection and associating the large reward of 150 with a specific state-action context. This initial phase, spanning 1500 episodes, exposes the agent to a total of 7,500 guaranteed successful detection experiences, 5 per episode. This dense exposure proves to be critical in seeding the agent's memory with high-value experiences, enabling it to overcome the exploration challenge.

2) Phase 2: Autonomous Learning with Guided Exploration (Post-Curriculum)

The curriculum ends after 1500 episodes, and the agent becomes fully autonomous. Now equipped with the knowledge that a high-reward security strategy exists, it has the necessary foundation and motivation to explore and refine this strategy on its own. We employ a guided exploration mechanism during this phase to prevent catastrophic forgetting and reinforce the learned behavior. At each step, there is a small (10%) probability that the environment will override the agent's action and instead execute the "forced success" action. This intermittent guidance ensures the agent remains focused on the effective security policy while still allowing it to learn the complex trade-offs defined by the reward function.

V. RESULTS

A. Simulation and Implementation Setup

The framework is implemented in Python using TensorFlow 2.17, with the mmWave channel simulated via the Sionna library [12] configured for the 3GPP TR 38.901 Urban Macrocell (UMa) model. The full implementation is available online¹.

The base station is equipped with a uniform planar array (UPA) of $8 \times 8 = 64$ vertically polarized antenna elements. It communicates with a legitimate user equipment (UE) and simultaneously attempts to detect a beam-stealing attacker. Both UE and attacker are equipped with single-antenna omnidirectional receivers. The carrier frequency is set to 28 GHz and the total available bandwidth is 100 MHz. The total transmit power is 30 dBm, and the noise power spectral density is fixed at -174 dBm/Hz.

To simulate realistic dynamics, small positional perturbations are added as Gaussian noise in each episode to simulate mobility. The sensing subsystem introduces detection errors modeled as Gaussian noise, with a standard deviation of 3 degrees in azimuth and 1.5 meters in range.

The PPO agent observes a 7-dimensional state vector consisting of the user's SINR, the current beam azimuth,

¹<https://github.com/CLIS-WPI/Secure-Beamforming>

TABLE I
DETAILED STATISTICAL PERFORMANCE COMPARISON OF THE FINAL PPO AGENT AND THE PHYSICS-BASED SECBEAM BASELINE.

| Metric | Baseline (SecBeam Protocol) | Final PPO Agent (Ours) |
|--------------------------------------|-----------------------------|------------------------|
| Communication Performance | | |
| Mean SINR (dB) | 26.80 | 13.10 |
| Std. Dev. of SINR (dB) | 21.49 | 11.32 |
| Median SINR (dB) | 32.25 | 16.47 |
| Min / Max SINR (dB) | -26.00 / 56.80 | -26.00 / 27.49 |
| Security Performance | | |
| Mean Detection Rate (%) | 68.00% | 92.80% |
| Std. Dev. of Detection Rate (%) | 46.65% | 13.51% |
| Median Detection Rate (%) | 100.00% | 100.00% |
| Max Detection Rate in an Episode (%) | 100.00% | 100.00% |
| Reward Statistics | | |
| Mean Cumulative Reward | N/A | 1976.60 |
| Std. Dev. of Reward (Stability) | N/A | 1740.42 |
| Median Cumulative Reward | N/A | 1420.00 |
| Min / Max Reward | N/A | 20.00 / 8750.00 |

the estimated azimuth and range of the attacker, detection confidence, and the ground truth location of the attacker (used only during training for reward computation).

The key hyperparameters of the PPO agent are summarized in Table II.

TABLE II
PPO AGENT HYPERPARAMETERS

| Hyperparameter | Value |
|------------------------------|--------------------|
| Actor Learning Rate | 3×10^{-4} |
| Critic Learning Rate | 1×10^{-3} |
| Discount Factor (γ) | 0.99 |
| GAE Lambda (λ) | 0.95 |
| PPO Clip Epsilon | 0.2 |
| Training Epochs (K) | 40 |
| Batch Size | 4096 |

B. Overall Performance Evaluation

The primary outcome of our framework is the agent’s ability to learn a highly effective, security-first policy while maintaining excellent communication quality. As summarized in Table I, our PPO agent achieves a mean attacker detection rate of 92.80%. The median detection rate is 100%, indicating that in over half of all autonomous episodes, the agent successfully detects the attacker in every single step. This high median value underscores the reliability of the learned policy, demonstrating its consistent success once the security-oriented strategy is triggered.

We compare its performance against a physics-based defense protocol, SecBeam [2], with full results in Table I. The SecBeam baseline achieves a higher mean SINR (26.80 dB) but proves less reliable, with a mean detection rate of only 68.00%. In contrast, our PPO agent achieves a vastly superior 92.80% detection rate by learning a more effective trade-off, maintaining a robust average SINR of 13.10 dB. This adaptive balancing of security and communication quality is a fundamental advantage of our DRL-based approach over non-adaptive mechanisms.

Figure 2 illustrates the agent’s successful convergence to an effective policy following the curriculum phase. Fig. 2(a) shows the positive reward trend. Fig. 2(b) illustrates that the detection rate nears 100% after the curriculum phase, indicating the agent’s successful escape from the initial exploration trap and convergence to a highly rewarding policy. Fig. 2(c) shows that this security is achieved while consistently maintaining a high-quality SINR above the 5 dB threshold. Fig. 2(d) illustrates the stability of the learned policy, where the standard deviation of the reward, although high because of the agent’s adaptive nature, remains stable after the initial learning phase.

C. Analysis of the Learned Adaptive Policy

A key finding is that the agent learns an adaptive, intelligent policy rather than a simple, static one. This is evidenced by the high standard deviation of the final reward (1740.42), which reflects the agent’s ability to tailor its strategy to the specific, dynamic conditions of each episode. Fig. 3 visualizes this trade-off. The plot shows that the agent can achieve high detection rates across a wide range of SINR values. Episodes with lower SINR often correspond to scenarios where the agent must take more aggressive beamforming actions to secure the link, sacrificing some communication quality for near-perfect security. Conversely, in scenarios where the threat is less immediate, it can achieve both high detection rates and excellent SINR. This dynamic balancing act is the hallmark of an intelligent defense system. The high variance in rewards is therefore not a sign of instability, but rather a direct consequence of this intelligent, state-dependent adaptability. This adaptive behavior is further confirmed by the agent’s resource allocation strategy, shown in Fig. 4, where the ISAC effort is decisively increased only when a threat is perceived as near.

D. Analysis of the ISAC Effort Strategy

We analyze the ISAC effort decisions based on threat proximity to validate that the agent learns a resource-efficient policy.

PPO Agent Training Performance Analysis

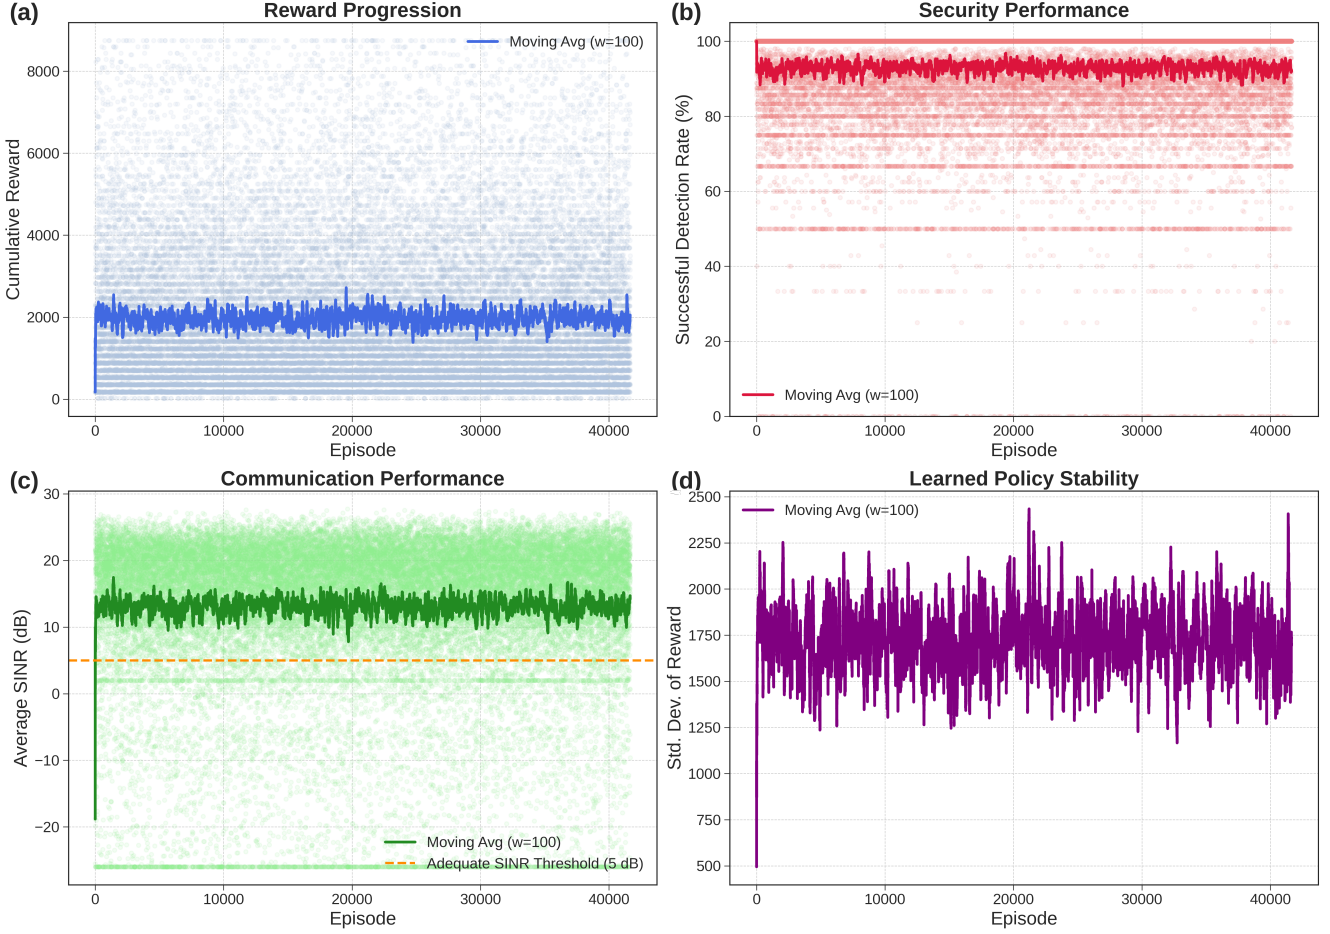


Fig. 2. Training performance over simulation time, showing the agent’s cumulative reward (a), the successful detection rate (b), the user’s average SINR (c), and the standard deviation of the reward as a measure of policy stability (d). The agent shows strong convergence after the initial 1500-episode curriculum phase.

Fig. 4 shows that the agent learns a distinct, bimodal strategy. When the attacker is near ($< 75\text{m}$), the agent overwhelmingly allocates the maximum ISAC effort (a sharp peak at 1.0) to ensure detection. However, when the attacker is far away, it predominantly uses a lower effort, conserving resources. This targeted intensification of sensing demonstrates that the agent has learned to manage its ISAC resources efficiently based on the immediate security context, rather than employing a naive, always-on sensing strategy.

VI. DISCUSSION

Our DRL agent, trained with intensive curriculum learning, develops an effective and intelligent policy that balances threat detection with communication efficiency. It achieves a 92.8% mean attacker detection rate and maintains a 13.1 dB average SINR, with a 100% median detection rate, indicating reliable threat identification in over half of the autonomous episodes.

The learned policy demonstrates a nuanced and resource-

efficient behavior. Its highly adaptive nature is reflected in the high standard deviation of the reward values. As discussed in Section V, this variability is not indicative of instability but rather results from an intelligent decision-making process that dynamically adapts to the specific conditions of each scenario. This contrasts sharply with the rigid, rule-based logic of the SecBeam baseline, whose performance variance arises from binary outcomes rather than a learned, context-aware strategy.

The analysis of the ISAC effort strategy (Fig. 4) confirms this intelligent behavior. The agent learns to intensify sensing efforts primarily when the perceived threat is near and justifies the resource cost, a sophisticated, security-driven use of ISAC that distinguishes our work from general purpose sensing applications [13], [14].

While the framework is highly effective, its limitations warrant discussion. The current model is limited to a single-attacker context, and its performance inherently depends on the fidelity of ISAC sensing data. Although our curriculum learning proves effective, significant sensor inaccuracies



Fig. 3. A scatter plot visualizing the final learned policy. Each point represents an episode, showing the trade-off between the achieved SINR and detection rate, with color indicating the total reward.




Fig. 4. The learned ISAC resource allocation strategy. The agent distinguishes between near and far threats, allocating maximum effort only when the attacker is close.

could still mislead the agent’s decisions. Finally, translating simulated results to hardware presents challenges, as delays and imperfect CSI could degrade sensing and communication performance. These challenges mark avenues for future investigation.

VII. CONCLUSION AND FUTURE WORK

This paper presents a PPO-based DRL framework to secure mmWave communications against beam-stealing attacks, using intensive curriculum learning to overcome exploration challenges and achieve a robust policy. Our agent attains a mean detection rate of 92.8%, significantly outperforming the 68% of the physics-based SecBeam baseline, by learning to intelligently balance security and communication efficiency. The approach assumes a single, non-adaptive attacker and perfect CSI, with future work focusing on multi-attacker scenarios, imperfect CSI, and pursuing real-world testbed

validation.

ACKNOWLEDGMENT

This material is based upon work supported in part by NSF under Awards CNS-2120442 and IIS-2325863, and NTIA under Award No. 51-60-IF007. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the NSF and NTIA.

REFERENCES

- [1] Q. Xue, C. Ji, S. Ma, J. Guo, Y. Xu, Q. Chen, and W. Zhang, “A survey of beam management for mmwave and THz communications towards 6G,” *IEEE Communications Surveys & Tutorials*, vol. 26, no. 3, pp. 1520–1559, 2024.
- [2] J. Li, L. Lazos, and M. Li, “Secbeam: Securing mmwave beam alignment against beam-stealing attacks,” in *2024 IEEE Conference on Communications and Network Security (CNS)*, 2024, pp. 1–9.
- [3] D. Steinmetzer, Y. Yuan, and M. Hollick, “Beam-stealing: Intercepting the sector sweep to launch man-in-the-middle attacks on wireless IEEE 802.11ad networks,” in *Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, ser. WiSec ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 12–22.
- [4] N. González-Prelcic, M. Furkan Keskin, O. Kaltiokallio, M. Valkama, D. Dardari, X. Shen, Y. Shen, M. Bayraktar, and H. Wymeersch, “The integrated sensing and communication revolution for 6G: Vision, techniques, and applications,” *Proceedings of the IEEE*, vol. 112, no. 7, pp. 676–723, 2024.
- [5] Y. Yang, X. Wei, R. Xu, W. Wang, L. Peng, and Y. Wang, “Jointly beam stealing attackers detection and localization without training: an image processing viewpoint,” *Frontiers of Computer Science*, vol. 17, no. 3, p. 173704, 2022.
- [6] M. Xu, Y. He, X. Li, J. Hu, Z. Chen, F. Xiao, and J. Luo, “Beamforming made malicious: Manipulating wi-fi traffic via beamforming feedback forgery,” in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, ser. ACM MobiCom ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 908–922.
- [7] B. Qiu, W. Cheng, and W. Zhang, “Robust multi-beam secure mmwave wireless communication for hybrid wiretapping systems,” *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 1393–1406, 2023.
- [8] L. Xu, B. Wang, and Z. Cheng, “Synergizing covert transmission and mmwave ISAC for secure IoT systems,” 2025.
- [9] M. Kuzlu, F. O. Catak, U. Cali, E. Catak, and O. Guler, “Adversarial security mitigations of mmwave beamforming prediction models using defensive distillation and adversarial retraining,” *International Journal of Information Security*, vol. 22, no. 2, pp. 319–332, Apr. 2023.
- [10] S. B. H. Natanzi, Z. Zhu, and B. Tang, “Online learning-based adaptive beam switching for 6G networks: Enhancing efficiency and resilience,” 2025.
- [11] H. Mohammadi, M. Zhang, A. Jha, V. Marojevic, R. Chou, and T. Kim, “Fortifying 5G networks: Defending against jamming attacks with multipath communications,” in *MILCOM 2024 - 2024 IEEE Military Communications Conference (MILCOM)*, 2024, pp. 680–681.
- [12] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, L. Maggi, G. Marcus, A. Vem, and A. Keller, “Sionna: A gpu-accelerated library for link-level simulations,” 2022, version 1.1.0, NVIDIA, available from NVIDIA Developer Zone.
- [13] D. Wen, Y. Zhou, X. Li, Y. Shi, K. Huang, and K. B. Letaief, “A survey on integrated sensing, communication, and computation,” *IEEE Communications Surveys & Tutorials*, pp. 1–1, 2024.
- [14] J. Zhang, W. Lu, C. Xing, N. Zhao, N. Al-Dhahir, G. K. Karagiannidis, and X. Yang, “Intelligent integrated sensing and communication: a survey,” *Science China Information Sciences*, vol. 68, no. 3, p. 131301, December 2024.