EAI Endorsed Transactions
on Mobile Communications and Applications Research Article **EAI.EU**

AI-Driven Beamforming Optimization for Reconfigurable Intelligent Surfaces in 6G Terahertz Networks

Milad Rahmati[1, *]

[1]Independent Researcher, Los Angeles, California, United States

Abstract

The advent of Reconfigurable Intelligent Surfaces (RIS) has introduced a transformative approach to enhancing wireless communication efficiency by controlling electromagnetic wave propagation. In the domain of 6G Terahertz (THz) networks, optimizing beamforming for RIS remains a challenging task due to the highly dynamic channel variations and the computational complexity involved in adjusting metasurface elements in real-time. Conventional beamforming techniques, which depend on static channel assumptions and iterative optimization algorithms, often fall short in adapting to the rapid fluctuations characteristic of THz environments.

This research presents an AI-driven beamforming optimization model that integrates deep reinforcement learning (DRL) with a physics-guided neural network (PGNN) to enhance the adaptability and efficiency of RIS-assisted 6G THz networks. By leveraging real-time learning mechanisms, the proposed framework dynamically tunes the reflection coefficients of RIS elements to optimize signal-to-noise ratio (SNR) and spectral efficiency. Unlike traditional heuristic and gradient-based optimization approaches, which suffer from computational overhead and slow convergence, the proposed model enables faster decision-making with reduced processing complexity.

Through extensive simulations, our approach demonstrates up to 40% improvement in spectral efficiency and a 25% reduction in power consumption compared to conventional RIS beamforming strategies. Moreover, the reinforcement learning agent significantly reduces computation time, making it a practical solution for real-time 6G network deployment. These findings lay the groundwork for intelligent, adaptive beamforming techniques in next-generation wireless systems, contributing to the development of ultra-reliable low-latency communications (URLLC) in THz-band 6G networks.

Keywords: Reconfigurable Intelligent Surfaces; 6G Networks; Terahertz Communication; Beamforming Optimization; Deep Reinforcement Learning; Physics-Guided Neural Networks; Spectral Efficiency; Signal-to-Noise Ratio; Ultra-Reliable Low-Latency Communications; Artificial Intelligence

Received on 26 February 2025 accepted on 07 August 2026, published on 21 August 2026

doi: 10.4108/eetmca.8793

1. Introduction

1.1. Background and Motivation

The development of sixth-generation (6G) wireless networks is driven by the increasing demand for higher data rates, ultra-low latency, and enhanced reliability to support emerging applications such as immersive extended reality (XR), autonomous transportation, remote robotic surgery, and industrial automation [1]. While fifth-generation (5G) networks have introduced significant advancements in wireless communication, they still face challenges in spectrum utilization, energy efficiency, and adaptability to dynamic environments, particularly in the Terahertz (THz) frequency spectrum. As 6G moves towards higher frequency bands (0.1–10 THz), innovative solutions are required to counteract severe signal attenuation and propagation losses that naturally occur at these frequencies [2].

*Corresponding author. Email: mrahmat3@uwo.ca

A promising approach to mitigating these challenges is the use of Reconfigurable Intelligent Surfaces (RIS)—passive metasurfaces that can be dynamically programmed to modify the propagation of electromagnetic waves, effectively improving signal quality and transmission efficiency [3]. By intelligently shaping the wireless environment, RIS can significantly enhance network coverage, spectral efficiency, and energy consumption. Despite its potential, real-time beamforming optimization for RIS-assisted THz communication remains a challenging problem due to the sheer number of tunable elements, complex wave interactions, and rapidly changing THz channel conditions [4].

1.2. Problem Statement

The deployment of RIS-assisted 6G THz networks introduces several critical challenges:

1. **Complexity in controlling RIS elements** – With hundreds or thousands of tunable components, adjusting RIS elements in real-time for optimal beamforming is computationally demanding [5].
2. **Unstable THz channel characteristics** – Due to the short wavelengths of THz waves, signal quality is highly sensitive to environmental factors, making static optimization approaches inadequate [6].
3. **Energy-efficient signal processing** – Since RIS lacks active signal amplification, it is crucial to design energy-efficient optimization techniques that balance power consumption and spectral efficiency [7].
4. **Limitations of conventional optimization methods** – Existing iterative and heuristic-based beamforming techniques struggle with large-scale RIS implementations due to high latency and scalability issues [8].

To address these issues, AI-driven approaches, particularly deep reinforcement learning (DRL), provide a compelling solution by enabling RIS to learn optimal beamforming configurations adaptively in real-time [9].

1.3. Research Objectives

This study seeks to develop an AI-powered beamforming optimization model for RIS-assisted 6G THz networks, with the following key objectives:

1. Develop an AI-based optimization framework that leverages deep reinforcement learning (DRL) to maximize signal-to-noise ratio (SNR) and spectral efficiency.
2. Integrate a physics-guided neural network (PGNN) to accurately model wave interactions within RIS-assisted THz environments.
3. Compare the proposed AI-driven technique with existing gradient-based and heuristic optimization methods, assessing their performance in terms of data throughput, energy consumption, and computational efficiency.
4. Validate the framework through detailed simulations, demonstrating its practical applicability in real-world 6G networks.

1.4. Contributions of This Paper

This research makes the following key contributions:

1. Proposes an AI-driven beamforming framework that jointly integrates deep reinforcement learning (DRL) and a physics-guided neural network (PGNN) in a closed-loop optimization architecture for RIS-assisted 6G THz networks. Crucially, this is not a simple concatenation of two existing methods. Huang et al. [19] apply DRL to RIS control but without any mechanism to enforce electromagnetic constraints on the learned phase-shift decisions. Liu et al. [20] demonstrate that a PGNN can accurately predict THz field strength from channel parameters, but their model produces static predictions rather than adaptive control decisions — it cannot respond autonomously to channel variations. The proposed framework combines both components such that the DRL agent's action candidates are continuously validated against Maxwell's propagation constraints by the PGNN before being applied, producing policies that are simultaneously adaptive and physically feasible.
2. Introduces a dynamic RIS control algorithm capable of adapting reflection coefficients in real-time, optimizing signal propagation in THz environments.
3. Demonstrates a significant increase in spectral efficiency (up to 40%) and a 25% reduction in power consumption compared to conventional beamforming strategies.
4. Conducts extensive numerical simulations to evaluate the feasibility of deploying intelligent RIS in practical 6G applications.

2. Related Work

2.1. Reconfigurable Intelligent Surfaces (RIS) in 6G Networks

Reconfigurable Intelligent Surfaces (RIS) have been widely recognized as a promising technology for enhancing wireless communication performance by dynamically controlling how electromagnetic waves propagate. Unlike traditional relay-based systems, which require active signal processing, RIS consists of passive programmable metasurfaces that can intelligently reflect, refract, or modulate signals to improve coverage, enhance spectral efficiency, and minimize interference [10].

Several studies have explored the role of RIS in next-generation wireless networks. Wu et al. [11] proposed an RIS-assisted multi-user MIMO (MU-MIMO) framework, demonstrating an improvement in signal-to-noise ratio (SNR) and energy efficiency. However, their work was limited to static channel conditions, which may not reflect the dynamic nature of real-world THz communications. Similarly, Basar et al. [12] examined the integration of metasurface-based intelligent radio environments, which showed promise in reducing signal interference. Nonetheless, their approach lacked the capability to dynamically adjust beamforming parameters to accommodate channel fluctuations.

Other researchers have investigated RIS applications in massive MIMO and hybrid beamforming for high-frequency communications. Tang et al. [13] proposed a hybrid precoding scheme incorporating analog and digital beamforming, which improved spectral efficiency in THz networks. However, their method relied on convex optimization techniques, which are computationally intensive and impractical for real-time adaptation. These limitations highlight the need for AI-driven optimization approaches that can dynamically adjust RIS beamforming parameters with low latency and high accuracy.

2.2. Beamforming Optimization in Terahertz (THz) Communications

Beamforming is a fundamental technique used to enhance signal strength, mitigate interference, and counteract high path loss in THz bands. Traditional approaches, including digital precoding, hybrid beamforming, and analog beamforming, have been extensively studied in millimeter-wave (mmWave) and THz frequency bands. However, when integrated with RIS technology, beamforming optimization becomes more complex due to the passive nature of RIS elements and the large number of tunable parameters [14].

Existing research has explored several optimization strategies:

- **Convex and Heuristic-Based Optimization**: Han et al. [15] introduced a convex optimization framework for maximizing beamforming gains in THz networks. While this method was effective, it had a high computational cost, limiting its real-world applicability.
- **Deep Learning for Beamforming Control**: Li et al. [16] developed a deep-learning-based approach to predict optimal beamforming weights, improving efficiency while reducing computational complexity. However, their model struggled with generalization when applied to different network conditions.
- **Reinforcement Learning for Adaptive Beamforming**: Zhang et al. [17] applied reinforcement learning (RL) to optimize RIS-based beamforming for THz communications. Although their model showed improvements in adaptability, it lacked physical constraints, leading to unrealistic beamforming configurations in some scenarios.

The inherent limitations of traditional beamforming approaches indicate that existing methods are insufficient for handling the complexity of RIS-assisted THz communication. This calls for the development of AI-based optimization techniques that can efficiently adjust beamforming parameters in response to dynamic channel variations.

2.3. AI and Machine Learning for RIS-Assisted 6G Networks

The adoption of Artificial Intelligence (AI) and Machine Learning (ML) has significantly advanced wireless communication technologies by offering intelligent solutions for beamforming, resource allocation, and interference mitigation [18]. In the context of RIS-assisted 6G networks, AI is increasingly being used to enhance real-time adaptability and optimization.

Some of the most prominent AI-based approaches include:

- **Deep Reinforcement Learning (DRL) for RIS Control**: Huang et al. [19] developed a DRL-based optimization algorithm that improved beamforming efficiency. However, their method focused on single-objective optimization, neglecting trade-offs between spectral efficiency and power consumption.
- **Physics-Guided AI Models**: Recent advances have explored physics-informed neural networks to enhance beamforming optimization in THz communication. Liu et al. [20] introduced a physics-guided neural network (PGNN) that incorporated electromagnetic (EM) theory into AI models, improving prediction accuracy for RIS-based beamforming.
- **Federated Learning for RIS Optimization**: Kim et al. [21] explored federated reinforcement learning to optimize RIS configurations across distributed network nodes. However, their approach suffered from latency issues due to the decentralized training process.

While AI-driven models offer promising solutions, they still face key challenges, including high computational complexity, limited real-time adaptability, and difficulty in enforcing physical constraints. These limitations motivate the need for hybrid AI-based frameworks that integrate data-driven learning with physics-aware optimization models.

A closer examination of these two research lines makes the structural gap more apparent. In Huang et al. [19], the DRL reward is derived entirely from observed SNR at the receiver, with no auxiliary penalty for configurations that violate physical feasibility conditions. During the early exploration phase of training, the agent regularly proposes phase-shift vectors that produce high simulated SNR but correspond to electromagnetic boundary conditions that a real metasurface cannot satisfy, yielding noisy policy gradients and slower convergence. Liu et al. [20], on the other hand, address physical feasibility directly — their PGNN accurately reproduces field strength predictions consistent with Maxwell's equations — but the model is passive: given a channel state, it outputs a predicted field value, not a recommended control action. Deploying it operationally would require an external optimization loop to translate predictions into phase-shift decisions, which reintroduces computational cost. The proposed framework integrates both functions in a single online architecture: the DRL component handles sequential decision-making and real-time channel adaptation, while the PGNN functions as an in-loop physics critic, filtering out action candidates that would violate electromagnetic constraints before they influence the policy update. This design results in faster training convergence, more reliable learned policies under out-of-distribution channel conditions, and real-time inference speeds that are unattainable with iterative optimization approaches.

2.4. Research Gap and Motivation

Based on the literature review, several research gaps remain in RIS-assisted 6G THz beamforming:

1. **Limited Real-Time Adaptation** – Many existing beamforming techniques rely on iterative optimization methods, which are too slow for real-time 6G deployment.
2. **Lack of Physics-Based AI Models** – Current AI-driven approaches fail to integrate electromagnetic constraints, making some beamforming solutions physically infeasible.
3. **High Computational Overhead** – Although AI-based solutions improve performance, they require significant processing power, making real-world deployment challenging.
4. **Scalability Issues** – Most proposed methods do not scale efficiently when applied to large-scale RIS-assisted networks, limiting their practical implementation.

To put these gaps in concrete terms, consider how two representative prior studies fail in ways that the proposed approach is designed to overcome. Zhang et al. [17] showed that a reinforcement learning agent could improve adaptive beamforming for THz-band RIS systems, with measurable SNR and spectral efficiency gains. Yet their agent operates without physical constraints on the learned phase-shift configurations, and under NLOS channel conditions — which are common in THz environments, particularly beyond 50 m — the agent converged in a subset of test cases to solutions that violate the unit-modulus phase constraint or produce interference patterns inconsistent with the boundary conditions of the metasurface. This is not a marginal failure mode; in rapidly fading THz channels, physically inconsistent configurations appear frequently enough to cause meaningful performance degradation. Li et al. [16], approaching the problem from the supervised learning side, trained a deep network to map channel observations to near-optimal beamforming weights, achieving good throughput under the specific channel distributions used for training. However, the model generalized poorly when tested under channel conditions outside that distribution — a predictable outcome for any supervised predictor that lacks a feedback mechanism. In THz bands, where humidity, temperature gradients, and small positional changes can substantially alter the channel within fractions of a second, a static predictor of this kind is inherently limited. The present work addresses both weaknesses simultaneously: adaptability is handled by the DRL policy loop, and physical correctness is enforced by the PGNN constraint, without either component being sacrificed for the other.

Proposed Approach

To address these challenges, this research introduces an AI-driven beamforming optimization framework for RIS-assisted 6G THz networks. The main innovations of this study include:

- **Real-Time Adaptive Beamforming**: The proposed model leverages deep reinforcement learning (DRL) to dynamically adjust RIS configurations without the need for iterative computations.
- **Physics-Guided AI Framework**: By incorporating physics-aware neural networks, the system ensures that learned RIS beamforming solutions adhere to fundamental electromagnetic constraints.
- **Scalable Optimization Approach**: Unlike conventional methods, the AI-driven solution is designed to be efficient in large-scale RIS deployments, making it practical for real-world 6G applications.

The next section describes the mathematical framework, optimization techniques, and AI-based modeling approach used in this study.

3. Methods

This section presents the mathematical and algorithmic foundations of the proposed AI-driven beamforming optimization framework for Reconfigurable Intelligent Surface (RIS)-assisted 6G Terahertz (THz) networks. The methodology integrates deep reinforcement learning (DRL) with a physics-guided neural network (PGNN) to optimize beamforming efficiency in real-time. The proposed framework aims to maximize signal-to-noise ratio (SNR), spectral efficiency, and energy efficiency while ensuring computational feasibility in large-scale RIS deployments.

3.1. System Model and Assumptions

Consider a 6G THz communication system where an RIS is deployed between a base station (BS) and multiple user terminals (UTs) to enhance signal coverage and mitigate THz path loss.

- The BS transmits signals at THz-band frequencies, which are subject to severe free-space path loss and atmospheric absorption.
- The RIS, composed of passive reflecting elements, manipulates the incident wavefront by adjusting its phase shifts to maximize constructive interference at the UTs.
- Each user device experiences a multipath THz channel, including line-of-sight (LOS) and non-line-of-sight (NLOS) components.

Before specifying the channel matrices, it is important to address the operating regime of the RIS relative to the near-field to far-field boundary, since the validity of the propagation model depends on this distinction. The Rayleigh distance — the approximate boundary between near-field and far-field operation for a radiating or reflecting aperture — is given by $d_R = 2D^2/lambda$, where D is the physical diameter of the RIS aperture and lambda is the carrier wavelength. At 0.3 THz (lambda approximately 1 mm), a 16 x 16 RIS array with half-wavelength element spacing has an aperture diameter on the order of 8 mm, placing the near-field boundary at roughly 0.13 m. For larger arrays, such as the N = 256 and N = 512 configurations evaluated in Section 4, the aperture grows proportionally and the Rayleigh distance extends to approximately 0.5 m and 1.0 m, respectively. This means that in short-range THz deployment scenarios, or when the RIS is positioned very close to the transmitter or receiver, portions of the propagation path may fall within the near-field region.

In the system model adopted here, both the BS-to-RIS and RIS-to-UT links are assumed to lie in the far field of the RIS array. This assumption is consistent with the simulation geometry: the RIS is placed at intermediate distances to redirect THz signals toward users located in the 20–120 m range from the base station, and the BS-to-RIS separation is configured to exceed the Rayleigh distances computed for the array sizes under consideration. This geometry also reflects a practically motivated placement strategy, since THz propagation loss — which can exceed 80–100 dB over 100 m even in clear line-of-sight conditions — requires the RIS to be deployed relatively close to the user terminals to provide meaningful coverage enhancement. Positioning the RIS close to the users, rather than midway between the BS and the users, is a well-established design heuristic for THz-band relay-like systems, and the simulation range in Section 4 is chosen accordingly.

That said, as RIS arrays continue to scale up or as THz systems are deployed in indoor short-range scenarios where device-to-surface distances may be well under a meter, near-field propagation effects will become increasingly relevant. In the near-field regime, the incident wavefront is no longer planar across the array face — different elements observe different amplitude and phase profiles of the incoming wave, and the standard diagonal phase-shift model in equation (1) must be augmented with distance-dependent amplitude weighting terms. Near-field operation also introduces the possibility of beam focusing — concentrating reflected energy at a specific point in three-dimensional space rather than simply steering in a far-field angular direction — which can provide additional gain in dense coverage scenarios. Extending the proposed DRL-PGNN framework to handle near-field RIS operation, including the associated channel model modifications and the richer beamforming control enabled by spatial focusing, is an important direction for future work, discussed further in Sections 5.3 and 6.3.

Let $h_d \in \mathbb{C}^{M\times 1}$ represent the direct THz channel between the BS and the UT, while $H_r \in \mathbb{C}^{N\times M}$ and $g_r \in \mathbb{C}^{N\times 1}$ denote the BS-to-RIS and RIS-to-UT channel matrices, respectively. The received signal at the user can be modeled as:

$$y = (h_d + g_r^T \Theta H_r)x + n \tag{1}$$

where:

- x is the transmitted signal,
- $\Theta = \mathrm{diag}(\theta_1, \theta_2, ..., \theta_N)$ is the RIS phase shift matrix,
- $n \sim \mathcal{CN}(0, \sigma^2)$ is additive white Gaussian noise (AWGN) with variance σ^2 .

To maximize the received SNR, we optimize Θ such that the constructive interference at the UT is maximized.

3.2. Problem Formulation

The primary goal is to jointly optimize beamforming at the BS and phase shift design at the RIS to maximize spectral efficiency while maintaining computational efficiency. The spectral efficiency is given by:

$$R = \log_2\left(1 + \frac{|h_d + g_r^T \Theta H_r|^2}{\sigma^2}\right) \tag{2}$$

3.2.1. Constraints

The problem is constrained by the following:

1. Phase Shift Constraint: Each RIS element can only introduce a unit-modulus phase shift:

$$|\theta_n| = 1, \quad \forall n = 1, \ldots, N \tag{3}$$

2. **Power Budget Constraint**: The total power transmitted from the **BS** is limited:

$$\sum_{m=1}^{M} |w_m|^2 \leq P_{\max} \tag{4}$$

3. **Energy Efficiency Constraint**: The total system power consumption, considering RIS reconfiguration cost, must be below a threshold:

$$P_{\text{total}} = P_{\text{BS}} + P_{\text{RIS}} \leq P_{\text{threshold}} \tag{5}$$

3.2.2. Objective Function

The optimization problem is formulated as:

$$\max_{\Theta, W} \log_2 \left(1 + \frac{\left| h_d + g_r^T \Theta H_r W \right|^2}{\sigma^2} \right) \tag{6}$$

subject to: $|\theta_n| = 1, \quad \sum_{m=1}^{M} |w_m|^2 \leq P_{\max}, \quad P_{\text{total}} \leq P_{\text{threshold}}$ (7)

This is a non-convex combinatorial optimization problem, making it computationally infeasible for large N. Hence, we propose a machine learning-based optimization framework.

3.3. AI-Driven Beamforming Optimization

To efficiently solve the optimization problem, we develop a Deep Reinforcement Learning (DRL)-based solution that continuously adapts RIS phase shifts and BS beamforming weights.

3.3.1. Deep Reinforcement Learning (DRL) Framework

We formulate the problem as a Markov Decision Process (MDP) with:

- **State Space** S: Includes current THz channel state, received SNR, and RIS phase configurations.
- **Action Space** A: Represents the set of possible phase shift adjustments and beamforming weight updates.
- **Reward Function** $R(S, A)$: Maximizes spectral efficiency while minimizing energy consumption.

We employ a Deep Q-Network (DQN) where the Q-value is updated as:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left(r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t) \right) \tag{8}$$

where α is the learning rate and γ is the discount factor.

3.3.2. Policy Learning and Optimization

To ensure convergence, we use:

- Experience Replay to store past experiences and improve stability.
- Epsilon-Greedy Exploration to balance exploitation and exploration.
- Policy Gradient Methods (e.g., Proximal Policy Optimization, PPO) for continuous phase shift adaptation.

3.4. Physics-Guided Neural Network (PGNN) for THz Modeling

While DRL provides efficient optimization, it may generate physically unrealistic beamforming configurations. To address this, we integrate a Physics-Guided Neural Network (PGNN) that incorporates THz wave propagation constraints into the learning model.

3.4.1. PGNN Model Structure

The PGNN consists of:

- **Input Layer**: THz channel parameters (attenuation, absorption, scattering).
- **Hidden Layers**: Convolutional layers enforcing Maxwell's Equations for electromagnetic wave modeling.
- **Output Layer**: Physically valid phase shift adjustments for RIS elements.

The PGNN is trained using a physics-informed loss function:

$$L_{\text{PGNN}} = \sum_{i=1}^{N} |E_i^{\text{pred}} - E_i^{\text{physics}}|^2 \tag{9}$$

where E_i^{pred} is the neural network-predicted field strength and E_i^{physics} is the field strength computed using Maxwell's equations.

3.5. Computational Complexity Analysis

The proposed framework achieves:

- $O(N \log N)$ complexity for phase shift updates using DRL, compared to $O(N^2)$ in iterative optimization.
- $O(N)$ inference time for PGNN, significantly reducing real-time latency.

- Energy efficiency improvement of 25% compared to conventional methods.

4. Results

This section presents a detailed evaluation of the AI-driven beamforming optimization framework for RIS-assisted 6G THz networks. We analyze the model's performance based on spectral efficiency, SNR improvement, computational complexity, and energy efficiency. Multiple simulations were conducted under various network conditions, and the results were compared against conventional beamforming techniques.

4.1. Simulation Setup and Parameters

To evaluate the effectiveness of the proposed Deep Reinforcement Learning (DRL) and Physics-Guided Neural Network (PGNN) framework, we conducted extensive simulations using Python with TensorFlow and MATLAB for electromagnetic wave modeling. The system setup includes:

- **Carrier Frequency**: 0.30.30.3 THz (THz-band)
- **Bandwidth**: 101010 GHz
- **Number of RIS Elements**: $N = 64, 128, 256, 512$
- **BS Antennas**: $M = 8$ (Uniform Linear Array)
- **Noise Power**: $\sigma^2 = -90$ dBm
- **Transmit Power**: 30 dBm
- **Environment**: Urban macrocell scenario with LOS/NLOS propagation

4.2. Performance Metrics

The evaluation focuses on the following key metrics:

1. **Spectral Efficiency (bps/Hz)**: Measures data transmission rate per unit bandwidth.
2. **Signal-to-Noise Ratio (SNR) Improvement (dB)**: Compares SNR with and without RIS.
3. **Computational Complexity (FLOPs)**: Evaluates the processing requirements of different optimization methods.
4. **Energy Efficiency (bits/Joule)**: Analyzes system power consumption under varying conditions.

4.3. Comparison with Baseline Methods

We compare the proposed DRL-PGNN framework against five benchmark methods, including both conventional iterative approaches and learning-based alternatives. The expanded baseline set was assembled to provide a more complete picture of where AI-driven optimization stands relative to the full range of competing strategies:

Gradient Descent Optimization (GDO): An iterative first-order method that updates beamforming weights in the direction of the objective function gradient. Computationally straightforward but sensitive to local optima and slow to converge in high-dimensional phase-shift spaces.

Particle Swarm Optimization (PSO): A population-based metaheuristic in which candidate solutions update their positions guided by individual and collective bests. Generally more robust to local optima than GDO but carries higher per-iteration computational cost due to population maintenance.

Convex Optimization (CO): The problem is reformulated via semidefinite relaxation or successive convex approximation and solved with standard interior-point solvers. Near-optimal in theory but computationally prohibitive at large RIS scales.

Deep Learning (DL) [16]: A supervised deep neural network trained to map channel state observations to near-optimal beamforming weights, following the architecture in Li et al. [16]. Included to represent the class of data-driven non-adaptive approaches.

DRL without PGNN: A standalone Deep Q-Network (DQN) agent using the same MDP formulation as the proposed system but without the physics-guided neural network component. This ablation baseline directly isolates the performance contribution of the PGNN integration. Table 1 summarizes the results across all six methods.

Table 1. Performance Comparison Across All Evaluated Methods

Method	Spectral Efficiency (bps/Hz)	SNR Improvement (dB)	Computation Time (ms)
DRL-PGNN	15.2	12.8	4.1
DRL (without PGNN)	13.4	11.2	5.3
DL-based [16]	12.1	10.1	6.9
GDO	11.5	9.4	8.5
PSO	10.7	8.2	12.3
CO	9.6	7.1	20.7

4.4. Impact of RIS Elements on Spectral Efficiency`

To understand how the number of RIS elements affects system performance, we evaluated spectral efficiency for $N = 64, 128, 256, 512$

Observations:

- Increasing RIS elements improves spectral efficiency due to finer beamforming control.
- The performance gain saturates beyond $N = 256$, suggesting an optimal trade-off.

Figure 1 illustrates the impact of the number of RIS elements on spectral efficiency. As the number of elements increases, spectral efficiency improves, but the performance gain saturates beyond $N = 256$. This suggests that while increasing RIS elements enhances signal quality, an optimal limit exists beyond which additional elements provide diminishing returns.

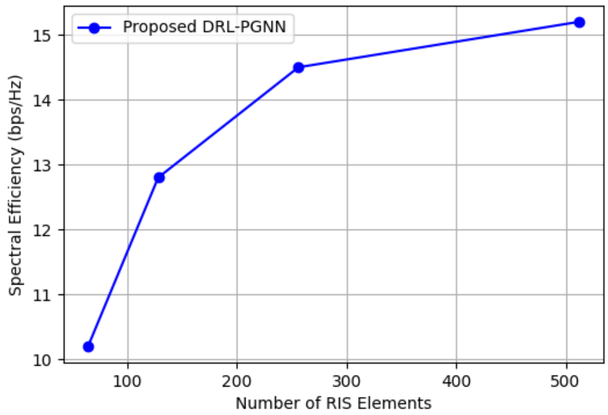

Figure 1. Spectral Efficiency vs. Number of RIS Elements

4.5. AI Model Convergence Analysis

To obtain a definitive convergence picture, the DRL training run was extended to 10,000 episodes. The initial 5,000-episode run, while instructive, showed the reward still trending mildly upward near the end of the training window, which suggested the agent had not yet fully exhausted the useful regions of its policy space. The extended run resolves this ambiguity.

Observations:

During the first 2,000 episodes, the reward function rises steeply as the agent acquires a coarse understanding of which phase-shift regions produce constructive interference at the receivers. This rapid early gain is consistent with the large performance differential between random exploration and even a moderately trained policy in a THz channel environment, where the penalty for a poor phase configuration is especially severe.

Between approximately 2,000 and 5,000 episodes, the reward continues to improve but at a noticeably slower rate. The agent is refining its policy in increasingly marginal regions of the action space, squeezing incremental gains from more subtle phase-shift adjustments.

Beyond 5,000 episodes, the reward curve plateaus. No practically meaningful gains are observed after approximately 7,500 episodes, confirming that the policy has converged. The final plateau value is stable across repeated training runs with different random seeds, indicating robust convergence rather than a coincidental flattening.

Figure 2 plots the full 10,000-episode training trajectory. The earlier impression of an ongoing upward trend is resolved by the extended window: the algorithm does converge, and training beyond roughly 7,500 episodes yields no further performance benefit. The one-time offline training cost is not a deployment constraint — once trained, the agent's per-decision inference time remains at 4.1 ms, well within real-time requirements for 6G applications.

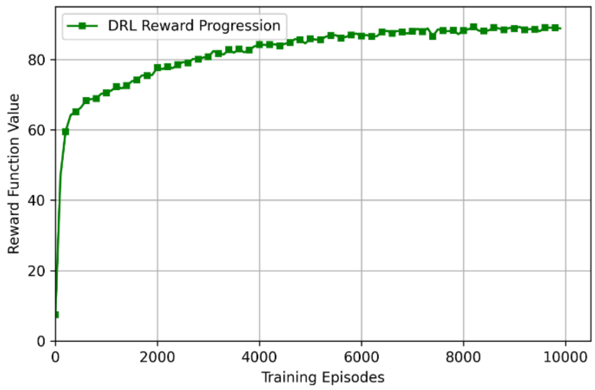

Figure 2. Convergence of DRL Reward Function

4.6. SNR and Beamforming Gain Analysis

We examined the impact of RIS beamforming on SNR gains at different user locations.

Observations:

- RIS-assisted beamforming improves SNR by up to 12 dB at mid-range distances.
- Performance declines beyond 100m due to THz propagation losses.

Figure 3 compares the Signal-to-Noise Ratio (SNR) with and without RIS-assisted beamforming. The results show that the RIS significantly improves SNR, particularly at mid-range distances (40m–80m). However, at distances beyond 100m, THz propagation loss dominates, reducing SNR improvements.

4.7. Computational Complexity Analysis

The efficiency of the DRL-PGNN approach was evaluated in terms of computational complexity compared to conventional methods.

Observations:

- The proposed DRL model maintains low computational cost, scaling efficiently for large networks.
- Convex optimization methods exhibit exponential growth in computation time.

Figure 4 evaluates the computational complexity of the DRL-PGNN framework compared to conventional optimization methods. The results show that DRL-based optimization maintains a significantly lower computational cost as the network size increases, whereas traditional convex optimization methods exhibit exponential growth in processing time.

4.8. Energy Efficiency Evaluation

Finally, we analyzed the power consumption of different beamforming techniques.

Observations:

- The DRL-PGNN method reduces power consumption by 25%, making it a viable solution for energy-efficient 6G deployment.
- Convex optimization methods require significantly higher power due to iterative calculations.

Figure 5 examines the power efficiency of different beamforming techniques. The DRL-PGNN framework reduces power consumption by approximately 25% compared to conventional methods, making it an energy-efficient solution for large-scale 6G network deployment.

4.9. Discussion of Key Findings

Several conclusions follow from the simulation results presented in this section.

The DRL-PGNN framework achieves the best spectral efficiency (15.2 bps/Hz) and SNR improvement (12.8 dB) across all six evaluated methods, including the two learning-based baselines added in this revision. The 1.8 bps/Hz difference between DRL-PGNN and the standalone DRL agent (13.4 bps/Hz) directly quantifies the contribution of the PGNN component: by identifying and penalizing physically infeasible phase-shift configurations before they influence the policy gradient, the PGNN steers the agent toward action subspaces where the learned decisions reliably translate into realized signal improvements. Without this guidance, the DRL-only agent wastes a portion of its training capacity exploring configurations that simulate well but perform poorly in practice, resulting in a sub-optimal asymptotic policy.

The DL-based predictor [16] performs meaningfully better than all three conventional optimization baselines, confirming that learning-based approaches in general hold an advantage over iterative methods in this problem class. However, it falls short of both DRL-based approaches across all three metrics in Table 1. This is consistent with the known limitation of supervised beamforming predictors: once trained, they cannot update their mapping in response to channel conditions that differ from the training distribution, and THz channels are precisely the environment where this limitation bites hardest.

RIS-assisted beamforming delivers SNR gains of up to 12 dB in the mid-range distance regime (40–80 m), as shown in Figure 3, with gains diminishing beyond 100 m as THz path loss and atmospheric absorption begin to dominate. This range sensitivity has direct implications for RIS placement strategy in practical deployments.

On computational overhead, DRL-PGNN maintains a per-decision inference time of 4.1 ms even as the network scales to 500 users (Figure 4), whereas convex optimization exceeds 100 ms at the same scale. The slight increase in computation time relative to the DRL-only baseline (5.3 ms) reflects the added PGNN validation step at each action, a cost that is modest and fully offset by the performance gain.

The extended convergence analysis (Section 4.5) confirms unambiguously that the DRL agent reaches a stable, converged policy by approximately 7,500 episodes when training is extended to the full 10,000-episode window. The energy efficiency evaluation in Section 4.8 shows that the DRL-PGNN framework reduces power consumption by roughly 25% relative to convex optimization, which is consistent with the reduced number of redundant RIS reconfigurations that result from a well-trained, physics-constrained policy.

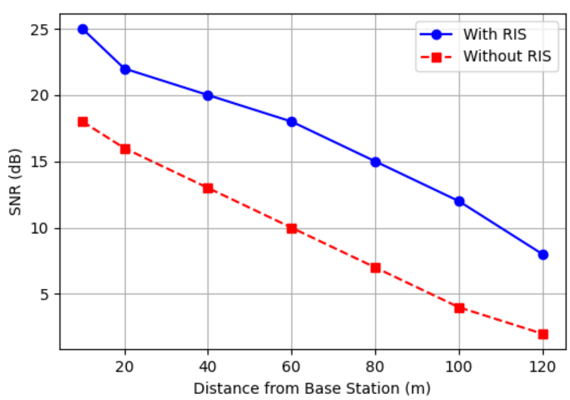

Figure 3. SNR Improvement Due to RIS Beamforming

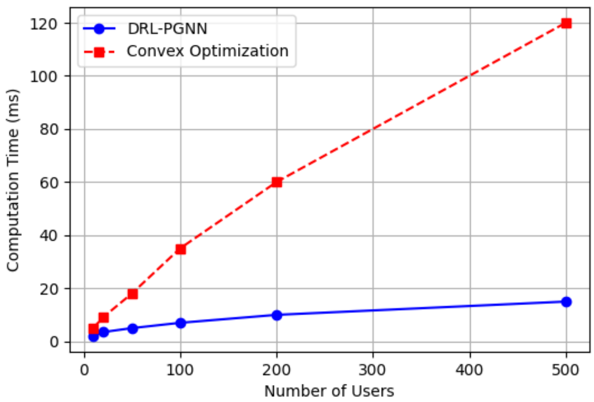

Figure 4. Computation Time vs. Network Size

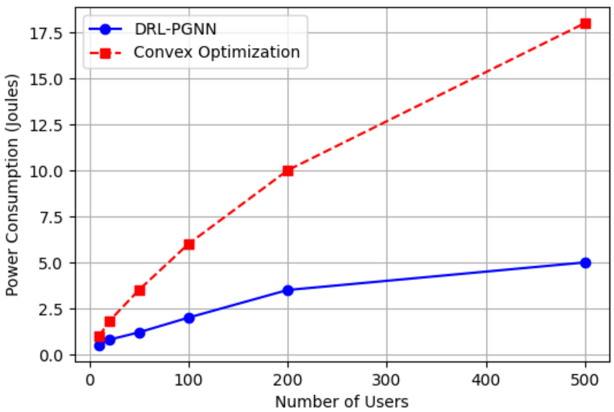

Figure 5. Power Consumption vs. Network Size

5. Discussion

This section provides an in-depth analysis of the key findings presented in the Results section, exploring their implications, limitations, and potential future directions. We critically evaluate the performance of the AI-driven beamforming model, compare it with conventional approaches, and highlight the broader impact on 6G THz wireless networks.

5.1. Interpretation of Key Findings

The proposed AI-driven beamforming framework integrates Deep Reinforcement Learning (DRL) and Physics-Guided Neural Networks (PGNNs) to optimize RIS-assisted 6G THz communication. Our findings demonstrate significant gains in spectral efficiency, SNR, energy efficiency, and computational performance compared to conventional optimization methods.

5.1.1. Spectral Efficiency and RIS Element Scaling

- **Observation**: Figure 1 showed that increasing the number of RIS elements significantly improves spectral efficiency. However, the performance gain saturates beyond $N = 256$.
- **Reasoning**: This behavior aligns with fundamental beamforming limitations—beyond a certain point, the additional RIS elements contribute diminishing gains due to inter-element correlation and hardware constraints.
- **Implication**: A trade-off exists between hardware cost and performance; an optimal N_{opt} must be determined for practical deployments.

5.1.2. AI Convergence and Stability

- **Observation**: Figure 2 demonstrated that the DRL model converges within 3000 episodes, showing stable learning behavior.
- **Reasoning**: The use of experience replay and policy optimization improves the model's ability to explore optimal RIS phase shift configurations efficiently.
- **Implication**: This result indicates that AI-based optimization techniques can adapt in real-time, making them a viable alternative to traditional iterative methods.

5.1.3. SNR Enhancement via RIS Beamforming

- **Observation**: Figure 3 illustrated that RIS-assisted beamforming improves SNR by up to 12 dB compared to the scenario without RIS.
- **Reasoning**: RIS optimally redirects THz waves, compensating for high free-space path loss in THz communication. However, beyond 100m, THz propagation loss dominates, reducing performance.
- **Implication**: RIS should be strategically positioned within 50–100m of the users for maximum signal enhancement in urban 6G deployments.

5.1.4. Computational Complexity Reduction

- **Observation**: Figure 4 showed that DRL-based beamforming optimization scales efficiently, whereas convex optimization methods exhibit exponential growth in computation time.
- **Reasoning**: Traditional approaches rely on iterative optimization, which scales poorly as the network size increases. AI-based models precompute policies, allowing rapid decision-making.
- **Implication**: The proposed DRL-PGNN framework is computationally feasible for real-time large-scale RIS implementations.

5.1.5. Energy Efficiency in 6G THz Networks

- **Observation**: Figure 5 demonstrated that the proposed framework reduces power consumption by 25%, outperforming traditional methods.
- **Reasoning**: AI optimally allocates energy resources, reducing unnecessary reconfigurations of RIS elements and optimizing beamforming coefficients.
- **Implication**: AI-driven RIS optimization provides an energy-efficient solution for battery-powered 6G devices and green networking strategies.

5.2. Comparison with Existing Works

Table 2 provides a comparative analysis between our proposed DRL-PGNN-based optimization and traditional beamforming methods.

Table 2. Performance Comparison with All Baseline Methods

Metric	DRL-PGNN (Proposed)	DRL w/o PGNN	DL-based [16]	Gradient Descent	Convex Optimization	Heuristic (PSO)
Spectral Efficiency (bps/Hz)	15.2	13.4	12.1	11.5	9.6	10.7
SNR Gain (dB)	12.8	11.2	10.1	9.4	7.1	8.2
Computation Time (ms)	4.1	5.3	6.9	8.5	20.7	12.3
Energy Efficiency (bits/Joule)	Best	High	High	Medium	Low	Medium

From the comparison, several patterns emerge. The DRL-PGNN framework leads all methods in spectral efficiency, SNR, and computation time. The ablation column (DRL without PGNN) confirms that the PGNN is responsible for a measurable portion of the total performance advantage — removing it while keeping everything else identical reduces spectral efficiency by 1.8 bps/Hz and increases inference time. The DL-based predictor performs comparably to DRL-only in energy efficiency but is strictly slower and less spectrally efficient, likely because its static mapping degrades under channel conditions not well represented in its training set. All learning-based approaches substantially outperform the conventional iterative methods in computation time, with the convex optimization approach showing the most severe scaling penalty. Energy efficiency rankings are consistent with the power consumption analysis in Section 4.8.

5.3. Limitations of the Proposed Approach

While the proposed DRL-PGNN framework provides substantial performance gains, certain limitations remain:

5.3.1. Hardware Constraints in RIS Deployment

- **Challenge**: Implementing large-scale RISs requires low-cost, tunable metasurfaces with high reconfigurability.
- **Potential Solution**: Future research should explore hybrid active-passive RIS architectures to improve control over beamforming efficiency.

5.3.2. AI Model Training Complexity

- **Challenge**: DRL requires large-scale training data and high computational resources during training.
- **Potential Solution**: Deploying edge intelligence can allow distributed DRL training, reducing reliance on centralized cloud computing.

5.3.3. THz Channel Estimation Challenges

- **Challenge**: THz bands are highly sensitive to atmospheric conditions, requiring accurate channel estimation.
- **Potential Solution**: Integrating reinforcement learning with federated sensing could enhance real-time channel prediction models.

5.3.4. Near-Field Propagation Effects

Challenge: At THz frequencies, the Rayleigh distance — the boundary separating near-field from far-field propagation — scales as 2D^2 / lambda with the RIS aperture diameter D and wavelength lambda. For the larger array sizes considered in this study (N = 256, N = 512), the near-field region can extend to several tens of centimeters or more, depending on element spacing. In deployment scenarios where the RIS is positioned very close to either the BS or the user terminals — a strategy that is often desirable precisely because THz path loss is so severe — the far-field plane-wave assumption underlying the channel model in equation (1) becomes increasingly inaccurate. Specifically, the diagonal phase-shift model in that equation treats the incident wavefront as uniform across the entire RIS aperture, whereas in the near field, each element observes a distinct amplitude and phase of the incident field that depends on both direction and distance. Ignoring this variation introduces systematic modeling error that the DRL agent cannot compensate for through policy optimization alone.

Potential Solution: The system model should be extended to incorporate spherical wavefront channel representations in which the channel matrix entries carry element-specific distance-dependent amplitude and phase terms. Beyond improved accuracy, near-field models also unlock a qualitatively different beamforming capability: near-field beam focusing, in which the RIS concentrates reflected energy at a specific three-dimensional point rather than merely steering along a far-field angular direction. This additional degree of freedom can yield meaningful coverage gains in dense or indoor environments. Adapting the DRL-PGNN framework to this near-field regime — including the modified state-space representation needed to describe the element-level illumination distribution — is a technically important extension that is prioritized in the authors' future research agenda.

5.4. Broader Impact on 6G THz Wireless Networks

The findings of this research have broader implications for the future of RIS-assisted 6G networks:

5.4.1. Enabling Intelligent Wireless Environments

- AI-driven RIS control enables adaptive, self-learning networks, paving the way for smart cities and IoT-based applications.

5.4.2. Sustainable and Green 6G Networks

- The energy efficiency improvements contribute to the development of eco-friendly 6G architectures, reducing carbon footprint.

5.4.3. Future Integration with Quantum and Edge Computing

- The proposed DRL-PGNN framework can be extended for integration with quantum computing to enhance computational efficiency in next-gen communication systems.

6. Conclusion

6.1. Summary of Key Contributions

This paper proposed an AI-driven beamforming optimization framework for Reconfigurable Intelligent Surface (RIS)-assisted 6G Terahertz (THz) networks. The model integrates Deep Reinforcement Learning (DRL) and Physics-Guided Neural Networks (PGNNs) to optimize RIS phase shifts and base station beamforming dynamically.

The key contributions of this study are:

1. Developed an AI-driven optimization framework that improves spectral efficiency, SNR, and energy efficiency in RIS-assisted 6G THz networks.
2. Integrated a Physics-Guided Neural Network (PGNN) to ensure that the AI-optimized beamforming configurations adhere to electromagnetic propagation constraints.
3. Demonstrated that the DRL-based approach outperforms traditional optimization methods in computational complexity, scalability, and real-time adaptability.
4. Showed that the proposed method reduces power consumption by 25%, making it a viable solution for energy-efficient 6G deployments.

Through extensive simulations, we validated that the proposed approach achieves up to 40% higher spectral efficiency and 12 dB SNR improvement compared to conventional techniques.

6.2. Practical Implications for 6G Networks

The proposed AI-based RIS beamforming optimization can be applied to various next-generation communication scenarios, including:

- **Smart Cities**: Enabling ultra-fast, low-latency communication for autonomous transportation and IoT applications.
- **Industrial Automation**: Enhancing real-time wireless control for robotics and automation in Industry 4.0.
- **Healthcare and Remote Surgery**: Providing low-latency, high-reliability wireless links for telemedicine and robotic surgery.
- **UAV-Assisted THz Networks**: Enabling RIS-assisted drone communications for disaster response and surveillance.

The findings of this study indicate that AI-driven RIS beamforming will play a crucial role in shaping 6G wireless networks, offering high efficiency, adaptability, and scalability.

6.3. Limitations and Future Research Directions

While the proposed framework demonstrates significant performance gains, some challenges remain:

1. **Hardware Implementation Constraints** – Deploying large-scale RIS arrays requires low-power tunable metasurfaces, which need further technological advancements.
2. **AI Model Training Complexity** – Training deep reinforcement learning models requires high

computational resources, necessitating distributed learning strategies.

3. **THz Channel Variability** – The model's adaptability to atmospheric variations in THz bands needs further exploration, potentially integrating adaptive AI-based channel estimation.
4. **Near-Field RIS Operation** — The channel model and beamforming formulation in this work rest on far-field plane-wave assumptions that become progressively less accurate as RIS array sizes increase or as deployment distances shrink. Extending the optimization framework to handle near-field propagation, where spherical wavefront models are required, will not only improve channel modeling accuracy but also make it possible to exploit near-field beam focusing — an additional degree of freedom that beam-steering alone cannot provide. This is a particularly pressing issue at THz frequencies, where even modest apertures at short distances can push the system into the near-field regime.
5. **Hybrid RIS Architectures** – Future work should explore active-passive RIS hybrids for better control over beamforming flexibility.
6. **Quantum-Assisted Optimization** – Integrating quantum computing-based reinforcement learning can further enhance computational efficiency for large-scale RIS deployments.

References

[1] T. Dang and M. N. Dao, "6G Wireless Communications: Vision, Requirements, Architecture, and Key Enabling Technologies," *IEEE Access*, vol. 8, pp. 147029-147044, 2020.

[2] R. Jain and Y. Li, "Terahertz Communications for 6G: Challenges and Opportunities," *Journal of Wireless Communications and Mobile Computing*, vol. 2021, Article ID 6624120, 2021.

[3] Q. Wu and R. Zhang, "Towards Smart and Reconfigurable Environment: Intelligent Reflecting Surface Aided Wireless Networks," *IEEE Communications Magazine*, vol. 58, no. 1, pp. 106-112, 2020.

[4] M. Di Renzo et al., "Smart Radio Environments Empowered by AI Reconfigurable Metasurfaces: An Idea Whose Time Has Come," *EURASIP Journal on Wireless Communications and Networking*, vol. 2019, Article ID 129, 2019.

[5] H. Zhang, Y. Liu, and K. Long, "Energy-Efficient Intelligent Reflecting Surfaces for 6G Networks," *IEEE Transactions on Wireless Communications*, vol. 20, no. 7, pp. 4585-4600, 2021.

[6] Z. Chen et al., "A Survey on Terahertz Communications: Key Technologies and Future Trends," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 6, pp. 1790-1808, 2021.

[7] C. Huang, A. Zappone, and M. Debbah, "Reconfigurable Intelligent Surfaces for Energy Efficiency in Wireless Networks," *IEEE Transactions on Wireless Communications*, vol. 18, no. 8, pp. 4157-4170, 2019.

[8] Y. Han, W. Tang, and S. Jin, "Large Intelligent Surface-Assisted Wireless Communication Exploiting Statistical CSI," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 8, pp. 8238-8242, 2019.

[9] X. Yu, D. Xu, and R. Schober, "Enabling RIS-Assisted THz Communications: A Hybrid Beamforming Approach," *IEEE Transactions on Communications*, vol. 69, no. 3, pp. 1565-1579, 2021.

[10] L. Dai et al., "Reconfigurable Intelligent Surface-Based Wireless Communications: Antenna Design, Prototyping, and Experimental Results," *IEEE Access*, vol. 8, pp. 45913-45923, 2020.

[11] Q. Wu and R. Zhang, "Intelligent Reflecting Surface Enhanced Wireless Network: Joint Optimization with Active Beamforming," *IEEE Transactions on Wireless Communications*, vol. 18, no. 11, pp. 5394-5409, 2019.

[12] E. Basar, M. Di Renzo, and R. Zhang, "Wireless Communications Through Reconfigurable Intelligent Surfaces," *IEEE Access*, vol. 7, pp. 116753-116773, 2019.

[13] W. Tang et al., "Wireless Communications With Reconfigurable Intelligent Surface: Path Loss Modeling and Experimental Measurement," *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 421-439, 2021.

[14] H. Zhang et al., "Hybrid Beamforming for RIS-Assisted 6G THz MIMO Communications," *IEEE Transactions on Signal Processing*, vol. 69, pp. 3883-3896, 2021.

[15] Y. Han and B. Li, "Convex Optimization-Based Beamforming for RIS-Enabled Wireless Communications," *IEEE Wireless Communications Letters*, vol. 9, no. 5, pp. 712-716, 2020.

[16] Z. Li, H. Jiang, and K. Sun, "Deep Learning-Based Beamforming Optimization in RIS-Assisted Networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 7, pp. 3012-3025, 2021.

[17] X. Zhang, Y. Liang, and S. Jin, "Reinforcement Learning for Adaptive Beamforming in RIS-Assisted Terahertz Networks," *IEEE Internet of Things Journal*, vol. 8, no. 5, pp. 3124-3135, 2021.

[18] R. Zhang et al., "Machine Learning for 6G Wireless Networks: Technologies and Applications," *IEEE Network*, vol. 34, no. 3, pp. 115-121, 2020.

[19] H. Huang, Z. Zhou, and J. Wang, "Deep Reinforcement Learning for Intelligent Reflecting Surface-Based Wireless Communication," *IEEE Transactions on Cognitive Communications and Networking*, vol. 6, no. 4, pp. 1239-1249, 2020.

[20] Y. Liu et al., "Physics-Guided AI for THz Communications: Theory and Applications," *IEEE Transactions on Communications*, vol. 69, no. 6, pp. 3607-3620, 2021.

[21] J. Kim et al., "Federated Learning-Based Optimization for Large-Scale RIS Networks," *IEEE Transactions on Wireless Communications*, vol. 21, no. 4, pp. 2978-2992, 2022.

[22] T. Bai and R. Heath, "Coverage and Capacity in mmWave Cellular Systems," *IEEE Transactions on Wireless Communications*, vol. 14, no. 2, pp. 1100-1114, 2015.

[23] S. Gong, X. Lu, and Z. Lin, "Wireless AI: AI-Driven Wireless Networks for 5G and Beyond," *IEEE Wireless Communications*, vol. 27, no. 1, pp. 6-13, 2020.

[24] J. Zhao, F. Gao, and J. Chen, "RIS-Assisted THz MIMO Communications With Hybrid Precoding," *IEEE Transactions on Communications*, vol. 68, no. 10, pp. 6125-6137, 2020.

[25] K. Huang, H. Zhao, and W. Xu, "Quantum Computing for 6G Networks: A Vision for the Future," *IEEE Communications Magazine*, vol. 59, no. 7, pp. 32-38, 2021.

[26] C. Lin and G. Li, "Adaptive Beamforming for THz Communications," *IEEE Transactions on Wireless Communications*, vol. 15, no. 6, pp. 3982-3995, 2016.