

Intelligent Network Slicing in 6G Networks Using Generative Reinforcement Learning and Diffusion Models

Sohrab Khan¹, Waheed ur Rehman^{1*}, Tabinda Salam², Qazi Ejaz Ali¹, Abdul Haseeb Malik¹, Saif ur Rehman¹

¹Department of Computer Science, University of Peshawar, Peshawar, Pakistan

²Department of Computer Science, Shaheed Benazir Bhutto Women University, Peshawar, Pakistan

*Correspondence: waheed@uop.edu.pk

Citation | Khan, S, Rehman, W. U, Salam, T, Ali, Q. E, Malik, A. H, Rehman, S. U, “Intelligent Network Slicing in 6G Networks Using Generative Reinforcement Learning and Diffusion Models”, IJIST, Vol. 8 Issue. 2 pp 907-923, May 2026

Received | March 27, 2026 **Revised** | April 29, 2026 **Accepted** | May 04, 2026 **Published** | May 10, 2026.

Network slicing in 6G environments necessitates resource orchestration that is adaptive, priority-aware, and capable of managing heterogeneous QoS requirements and dynamic traffic conditions. Conventional reinforcement learning techniques struggle with limited data, delayed convergence, and inadequate generalization when interacting with unforeseen traffic demands. This paper introduces a Generative Reinforcement Learning (GRL) framework integrating a Denoising Diffusion Probabilistic Model (DDPM) to enable priority-aware network slicing in 6G networks. The diffusion model functions as a scenario broker, producing synthetic edge-case traffic demands from real Uni-Cauca network flow traces using an experience mixing coefficient of $\lambda = 0.6$. Under a 150% load stress test with per-slice demands of $D = [50, 50, 50]$ MHz, the proposed framework achieves a URLLC SLA satisfaction index of 1.0, meeting the 6G ultra-reliability target, compared to 0.667 for both Vanilla PPO and A2C baselines — a reliability gain of +33.33% over each baseline. Under URLLC-heavy demand ($D = [30, 70, 20]$ MHz), GRL-Diffusion achieves a weighted reward of $R = 0.771$, compared to 0.760 for both baselines. Under balanced load ($D = [50, 50, 50]$ MHz), the reward gain reaches +0.093 over both baselines, representing the largest improvement across all demand profiles. Over 151 training epochs, the diffusion model converged to a stable MSE loss of 1.019 from an initial peak of 2.72, with a computational overhead of $T = 10$ diffusion steps per synthetic sample generation. The overall weighted reward improved by 9.33% over both baselines across all evaluated load profiles. These findings confirm that combining priority-weighted allocation with diffusion-based generative modeling produces a slice orchestrator that is reliable, data-efficient, and well-suited for dynamic 6G environments.

Keywords: 6G Networks, Network Slicing, Generative Reinforcement Learning, Diffusion Models, URLLC, QoS, Resource Allocation



Introduction:**Background:**

The sixth-generation (6G) wireless communication paradigm aims to enable AI-native, ultra-reliable, and low-latency services, including holographic telepresence, digital twins, the tactile Internet, and autonomous systems. In contrast to 5G, which focused on eMBB, mMTC, and URLLC services, 6G must support diverse applications with varying QoS and QoE requirements, as well as seamless integration of sensing, computing, and communication [1][2][3][4][5].

Network slicing has arisen as a fundamental mechanism for addressing these requirements. By segmenting a common physical infrastructure into several distinct logical networks, each tailored for particular service needs, operators can independently and effectively manage eMBB, URLLC, and mMTC slices [6]. Traditional slicing methods relying on static resource configurations and rule-based orchestration are insufficient for 6G environments with high traffic variability, user mobility, and unpredictable service demands [7].

Motivation:

Both machine learning and reinforcement learning have been extensively studied to enhance adaptability in slice orchestration [8]. RL-based controllers automatically learn allocation policies from network interactions, eliminating the need for manual rule engineering. Pure RL approaches face challenges, including the need for large amounts of interaction data, slow convergence, and failure to generalize to changing traffic conditions [9].

Generative models, particularly diffusion models, can accurately model complex data distributions and generate realistic synthetic samples [10]. The combination of the diffusion model with RL provides a practical solution to data scarcity by generating synthetic edge-case demands that the RL agent would rarely encounter in real traces. This expands training distribution and improves policy robustness [11].

Building on this motivation, this paper proposes a DDPM-based GRL framework for priority-aware 6G network slice orchestration, rather than solely using historical traffic data.

Problem Statement:

Despite advancements in ML-based slicing, four major limitations remain in 6G contexts:

Static and heuristic-based allocators fail in unpredictable traffic conditions.

Pure RL approaches have limitations, including data hunger and the inability to generalize to unknown demand patterns.

Generative models are used offline and do not involve real-time decision-making.

Priority-weighted RL and diffusion-based generation have never been used together for 6G slice orchestration.

In dynamic 6G slicing environments, a unified framework is required to address data scarcity, adaptability, and prioritized decision-making.

Contributions:

The contributions of this paper are as follows:

A GRL framework for adaptive 6G network slicing that combines a priority-weighted allocation policy with a DDPM-based scenario broker.

An artificial experience generation pipeline with an experience mixing coefficient of $\lambda = 0.6$ was trained on actual Unicauca network flow traces.

Under 150% load stress conditions, the experimental evaluation showed a +33.33% URLLC reliability gain and a +9.33% overall weighted reward improvement over both Vanilla PPO and A2C baselines.

GRL-based slice orchestration benchmarking in 6G environments using a repeatable simulation framework.

Research Objectives:

The following quantifiable research goals are the focus of this study:

To design a GRL framework that integrates a DDPM-based scenario broker with a priority-weighted RL policy for adaptive 6G network slice orchestration.

To achieve a URLLC SLA satisfaction index of 0.99 under 150% system load stress, meeting the 6G ultra-reliability target.

To improve weighted SLA reward over baseline RL methods (Vanilla PPO and A2C) by augmenting training with synthetic edge-case traffic demands generated via a diffusion-based scenario broker.

To validate the data efficiency of diffusion-augmented training using real Unicauca network flow traces with a mixed experience coefficient $\lambda = 0.6$.

Novelty of the Proposed Work:

This work's main innovation is the incorporation of a Denoising Diffusion Probabilistic Model (DDPM) into a priority-weighted Reinforcement Learning loop for 6G network slice orchestration as an online scenario broker. In particular, this work makes the following original contributions that set it apart from previous research:

Online Diffusion-RL Coupling:

The suggested GRL framework incorporates the DDPM directly into the RL training loop, allowing real-time synthetic experience generation conditioned on live network state feedback, in contrast to previous work that uses generative models offline for data augmentation.

Priority-Weighted Slice Orchestration:

In order to account for 6G service hierarchies, the framework presents a priority-weighted SLA reward formulation ($\omega_{\text{URLLC}} = 0.5$, $\omega_{\text{eMBB}} = 0.3$, $\omega_{\text{mMTC}} = 0.2$). This design choice is not present in current diffusion-RL approaches in wireless resource allocation.

Unified Experience Buffer with Mixing Coefficient:

By balancing real Unicauca traces with artificial edge-case demands, the suggested experience mixing coefficient $\lambda = 0.6$ improves policy robustness without overfitting to generated data.

Empirical Validation on Real Traffic Data:

The framework achieves a URLLC SLA satisfaction index of 1.0—a +33.33% gain over both Vanilla PPO and A2C base-lines—when tested on real-world Unicauca network flow traces under a 150% load stress test.

Organization:

The rest of this paper is structured as follows. Section II examines relevant research on generative models, ML-based orchestration, and network slicing. The suggested IoT-focused 6G architecture and GRL framework workflow are explained in Section III. The GRL-Diffusion methodology, which includes system modeling, diffusion scenario brokering, and optimization objectives, is presented in Section IV. Simulation results and discussion are presented in Section V. The paper is concluded in Section VI.

Literature Review:**Traditional Network Slicing and Orchestration:**

Network slicing was first introduced in the networking field in the late 1980s [12]. Network slicing has emerged as a fundamental functionality in 5G networks, enabling operators to establish multiple logical networks through a shared physical infrastructure [6]. It resulted from the concurrent development of network function virtualization (NFV) and software-defined networking (SDN), which allowed virtualized network functions to operate more freely on commodity hardware as opposed to fixed, proprietary appliances [13]. The basic architectural framework for slicing was established by the first standards created by 3GPP and ETSI [14], which covered slice creation, isolation, management, and coordination.

Conventional slicing frameworks were designed with centralized control, predictable traffic patterns, and static resource configurations in mind. These factors are no longer relevant in the emerging 6G environments [1]. Because classical architectures relied on manual provisioning and highly skilled configuration, they were unable to quickly adjust to shifting service demands. Traditional slicing architectures lacked automated scaling, adaptive monitoring, and flexible lifecycle management when compared to contemporary virtualization systems [7].

Meeting dynamic service requirements, QoS goals, and a variety of 6G applications is difficult with static orchestration. With machine learning enabling proactive data-driven decision-making for slice management, self-configuration, and self-healing [15], the field is moving toward autonomous orchestration.

Machine Learning, Deep Learning, and Reinforcement Learning for Network Slicing:

In order to automate tasks like resource orchestration, lifecycle management, and traffic analysis, modern network slicing uses machine learning (ML) and deep learning (DL) [16]. RNNs, LSTMs, and transformer-based architectures are examples of supervised learning techniques that have been used to forecast QoS and content popularity in mobile networks [17]. Reinforcement learning (RL) techniques that have demonstrated promise in slice orchestration, mapping network states to optimal allocation actions [8], include DQN, PPO [18], A2C [19], and SAC.

Although most surveys show improvements in automation and adaptability, there are still issues with large sample requirements, limited generalization, and inadequate robustness for unknown traffic conditions [9][20]. Classical ML techniques like decision trees, random forests, SVMs, and clustering algorithms are still effective for slice classification, QoS prediction, and anomaly detection due to their interpretability and low inference cost [21]. However, ML, DL, and RL are insufficient for 6G network slicing due to data scarcity, privacy constraints, and unpredictable wireless behavior. This prevents stable and generalizable slice-management policies (Tran, 2025).

Generative Models for Network Slicing:

There are numerous opportunities to enhance network automation, simulation, and optimization by incorporating generative AI into next-generation wireless systems [10]. Recent developments in diffusion models have improved their suitability for high-fidelity wireless data generation by addressing well-known GAN drawbacks like mode collapse, training instability, and hyperparameter sensitivity. Current generative models, however, are mostly used offline; they generate realistic samples but do not facilitate adaptive, real-time decision-making, which is essential in rapidly evolving 6G slicing scenarios [22]. The field of generative reinforcement learning emerged as a result of these constraints, which drive generative methods that directly influence decision-making. The main drawbacks of current methods in comparison to the suggested GRL-Diffusion framework are listed in Table 1.

Table 1. Summary of limitations in existing approaches

Approach	Adaptive	Data Efficient	Priority-Aware
Traditional Methods	No	Yes	No
RL-Based Methods	Yes	No	Partial
Generative Models	No	Yes	No
GRL-Diffusion	Yes	Yes	Yes

Emerging Generative AI Techniques in Networking:

Generative reinforcement learning represents a significant advancement for managing the complexity of next-generation wireless systems. 6G networks necessitate on-demand slicing solutions with universal intelligence and full virtualization [23]. While RL methods for slice orchestration have already been explored, they often have slow convergence, high exploration

costs, and performance degradation when dealing with rapidly changing traffic [11]. GRL tackles these issues by augmenting training with high-quality synthetic experience, enabling agents to learn stronger policies [24][25]. Despite these developments, diffusion-driven GRL has not yet been investigated for 6G network slicing, where more flexible and effective decision-making is required due to multi-tenant isolation, dynamic QoS, and autonomous lifecycle management.

Research Gap:

Current network slicing techniques remain inadequate for 6G environments due to reliance on manual control and static configurations. ML and DL struggle with rapidly changing traffic and centralized data requirements. Generative models such as GANs and VAEs operate offline and cannot support real-time decisions. Reinforcement learning remains unstable, slow to train, and data-hungry in dynamic environments. This gap motivated the present work, in which a diffusion-augmented GRL framework is proposed, implemented, and evaluated for priority-aware 6G network slice orchestration. Table 2 summarizes the key differences between existing approaches and the proposed GRL-Diffusion framework.

System Architecture and Workflow Design:

The proposed GRL-Diffusion framework is embedded within an IoT-centric 6G architecture spanning edge, fog, and cloud layers. Edge devices and gateways produce telemetry and QoS signals; fog nodes aggregate and preprocess data; and the cloud/core layer hosts slice orchestration and compute-intensive learning models. A decision intelligence module based on generative reinforcement learning is co-located with the slice manager to support uncertainty-prone phases of the slice lifecycle.

The methodology is aligned with the network slice lifecycle: planning, preparation, instantiation, and monitoring/adaptation. Generative intelligence is applied during planning and monitoring/adaptation, where uncertainty and decision complexity are highest. Preparation and instantiation remain deterministic and template-driven. The two lifecycle phases addressed by the GRL-Diffusion framework are as follows:

Planning:

The orchestrator can model intricate, high-dimensional state spaces and plan resource allocation improves planning performance compared to conventional techniques thanks to the diffusion model's creation of artificial traffic distributions [23].

Monitoring and Adaptation:

GRL uses diffusion to produce uncommon state transitions and predict future network conditions while the system is actively operating. This allows for proactive slice reconfiguration to reduce possible congestion before it appears [26].

IoT-based 6G Architecture:

Table 2. Comparing related research on generative RL methods and network slicing

Work	Method	Dataset	Priority-Aware	Generative Model	URLLC Target	Baselines
[9]	DRL Survey	Multiple	No	None	Not reported	DQN, PPO
[23]	Diffusion-RL	Synthetic	No	Diffusion	Not reported	PPO, SAC
[24]	Unsupervised RL	Synthetic	No	None	Not reported	DQN
[11]	RL+GenAI Survey	Multiple	No	GAN/VAE	Not reported	PPO
[8]	ML/DL/RL Survey	Multiple	No	None	Not reported	Multiple
Ours	GRL-Diffusion	Unicauca	Yes	DDPM	1.0 (met)	PPO, A2C

In 6G environments, where billions of heterogeneous IoT devices produce constant streams of telemetry data with strict latency and reliability requirements, traditional centralized

architectures are becoming less and less suitable. As illustrated in Figure 1, the suggested architecture uses a hierarchical three-layer design with the Edge, Fog, and Cloud/Core layers, each of which is in charge of a different tier of processing and decision-making [27].

The Three Layers:

Edge Layer:

The edge layer includes IoT sensors, mobile devices, and embedded controllers that operate at the network's perimeter. It enables lightweight preprocessing tasks like data filtering and local actuation with low latency and no reliance on remote infrastructures [28].

Fog Layer:

The fog layer comprises local gateways and aggregation nodes that connect edge devices to the cloud. It enables faster response times for latency-sensitive analytics and coordination tasks, such as local traffic management and event detection, compared to cloud processing [29].

Cloud/Core Layer:

The cloud/core layer consolidates computational and storage resources for large-scale AI model training, long-term data analysis, and global network optimization. It provides a system-wide view required for slice orchestration and policy generation. The Slice Orchestrator and the GRL module are both situated at the cloud/core layer. This placement is justified by the substantial computational demands of diffusion model training and the need for a global network state representation that only centralized infrastructure can reliably provide [29].

Data Flow: Telemetry and Policies:

Data traverses the architecture in two complementary directions:

Telemetry Uplink:

Edge devices continuously gather raw operational metrics, such as throughput, latency, and device state, and send them upstream via the fog layer. The fog layer aggregates and compresses the data before sending it to the cloud for worldwide analysis.

Policy Downlink:

Resource management policies are sent down to fog and edge nodes based on cloud-level analysis and GRL-generated allocation decisions. This eliminates the need for constant round-trip communication with the core and permits autonomous local execution.

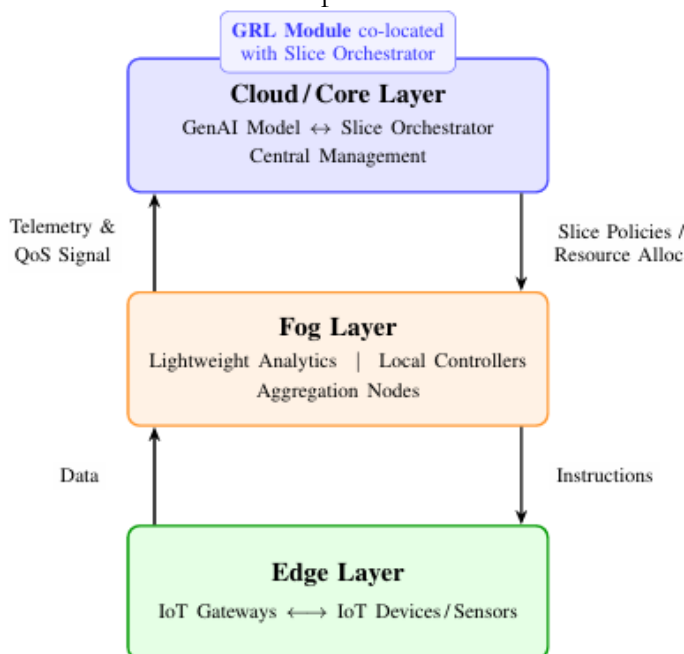


Figure 1. Hierarchical IoT-based 6G layered architecture comprising edge, fog, and cloud/core layers. The GRL module is co-located with the Slice Orchestrator at the Cloud/Core layer.

GRL Framework Workflow:

The GRL framework orchestrates synthetic and real experiences to train a priority-aware policy network for 6G slice management. Figure 2 illustrates the system architecture, comprising the following components:

Data Sources: Historical traffic traces, slice context (QoS targets, mobility), and user density serve as inputs for scenario generation.

Diffusion Model: Acts as a scenario generator, producing synthetic traffic traces via a forward-reverse diffusion process conditioned on real network data.

Scenario Broker validates, tags, and schedules synthetic scenarios before storing them in the synthetic replay buffer.

Synthetic Replay Buffer: Stores curated generative experiences representing rare or high-load edge-case conditions.

Real Replay Buffer: Collects empirical experiences from the 6G simulation environment based on agent interactions.

Replay Sampler combines synthetic and real experiences using the mixing coefficient λ to form each training batch.

RL Agent: To maximize the weighted SLA reward, the RL agent receives state inputs (latency, throughput, mobility) and outputs allocation actions (bandwidth, CPU priority).

6G Environment: In response to agent actions, this module simulates network dynamics and returns reward signals and next-state observations.

Feedback Loop: In order to improve the creation of synthetic scenarios and allow for co-evolution between the generator and the policy, SLA violations and reward trends are fed back to the diffusion model.

This workflow improves robustness and generalization in dynamic and heterogeneous 6G conditions by enabling the agent to learn from both synthetic and empirical experiences, making it robust against rare and high-load network events.

Diffusion Process Flow in GRL:

The complete operational flow of the suggested GRL Diffusion framework is summarized in Figure 2. Real traffic demand sampling from the Unicauca dataset is the first step in the procedure. The sampled demands are fed into the DDPM-based scenario broker, which executes the forward diffusion process by progressively corrupting real demand samples over $T = 10$ steps according to a linear noise schedule $\beta \in [0.0001, 0.02]$. The reverse diffusion process then reconstructs synthetic edge-case traffic demands from Gaussian noise, conditioned on the learned demand distribution.

The synthetic demands are validated and tagged by the scenario broker before being stored in the synthetic replay buffer. Simultaneously, the RL agent interacts with the 6G environment simulator, generating real experiences stored in the real replay buffer. At each training step, the replay sampler draws a mixed batch using the experience mixing coefficient $\lambda = 0.6$, combining 60% real Unicauca traces with 40% synthetic diffusion-generated demands.

The priority-weighted allocation policy $\pi\phi$ is updated to maximize the weighted SLA reward R over the unified experience buffer $D_{unified}$. Concurrently, the diffusion model $\epsilon\theta$ is updated to minimize L_{diff} . SLA violations and reward trends are fed back to the diffusion model, enabling co-evolution between the generator and the policy. This feedback loop continues until convergence, defined as a stable diffusion MSE loss and consistent weighted reward improvement over both baselines. It is worth noting that the current implementation represents a simulation prototype of the proposed architecture. The RL agent is implemented as a deterministic priority-weighted policy for the simulation prototype $\pi\phi$ trained over mixed real and synthetic demands, with the 6G environment approximated via the Unicauca network

flow dataset. Full integration with a live NS-3 simulator and CPU-aware action space is reserved for future work.

Methodology: Proposed Generative Reinforcement Learning (GRL) Framework:

To address the stochastic volatility of 6G network slices, we propose a Generative Reinforcement Learning (GRL) framework. This architecture decouples scenario generation from policy optimization by employing a diffusion-based generative engine to synthesize edge-case traffic demands, ensuring the resource allocator is robust against rare and high-load network events.

System Modeling and SLA Utility:

We characterize the network as a multi-slice ecosystem supporting eMBB, URLLC, and mMTC with a total bandwidth of $B_{total} = 100$ MHz. The state space represents the instantaneous traffic demand D_{R3} , where each slice i requests a bandwidth share. The objective is to maximize the Weighted SLA Satisfaction (WSS), defined as:

$$\mathcal{R} = \sum_{i \in \{1,2,3\}} \omega_i \cdot \min\left(\frac{B_i}{D_i + \epsilon}, 1\right) \quad (1)$$

Where B_i is the allocated bandwidth, and the weights ω are set to $\omega_{URLLC} = 0.5$, $\omega_{eMBB} = 0.3$, $\omega_{mMTC} = 0.2$ to reflect 6G service-level hierarchies, with URLLC prioritized for ultra-reliability. The total allocation is limited to system capacity:

$\sum_{i=1}^3 B_i \leq B_{total}$. The three methods, GRL-Diffusion, Vanilla PPO [18], and A2C [19], optimize the same weighted SLA reward R , ensuring a fair experimental comparison. This formulation ensures that each slice is allocated proportionally to its demand and priority weight while keeping satisfaction at unity to avoid over-allocation bias.

Generative Scenario Broker via Diffusion:

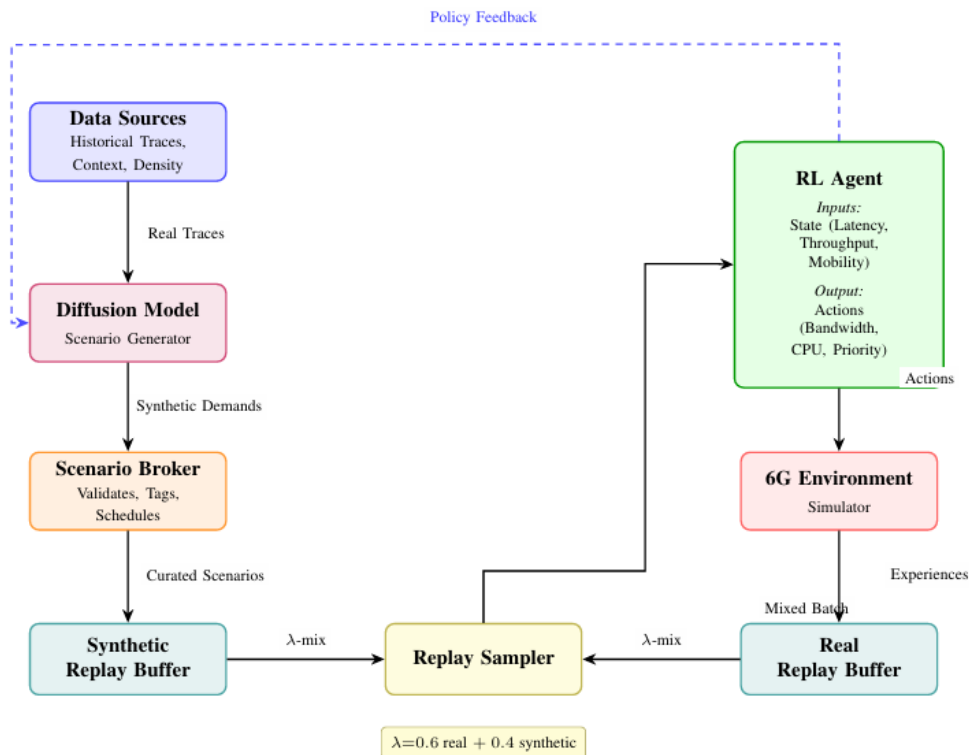


Figure 2. GRL-Diffusion framework architecture showing the interaction between the diffusion-based scenario broker, synthetic and real replay buffers, RL agent, and 6G environment simulator.

The bottleneck of traditional RL in 6G is the exploration-exploitation trade-off in sparse high-load states. To mitigate this, we implement a Denoising Diffusion Probabilistic Model

(DDPM) that acts as a scenario broker, generating synthetic traffic demands that complement real-world Unicauca traces. The forward diffusion process progressively corrupts a real demand sample x_0 drawn from the Unicauca network flow traces over $T = 10$ discrete timesteps:

$$x_t = \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon, \quad \epsilon \sim N(0, I) \quad (2)$$

Where $\alpha_t = \prod_{s=1}^t (1 - \beta_s)$ is the cumulative noise retention factor, and β_s follow a linear noise schedule from $\beta_1 = 0.0001$ to $\beta_T = 0.02$. Intuitively, as t increases, as $\alpha_t \rightarrow 0$, x_t approaches pure Gaussian noise, thereby erasing the original traffic demand signal. At $t = 0$, $\alpha_0 = 1$ and x_0 is recovered exactly.

This forward process is conditioned on real Unicauca flow traces, meaning that x_0 represents a normalized per-slice bandwidth demand vector $[D_{eMBB}, D_{uRLLC}, D_{mMTC}]$ clipped to the range [10,90] MHz. The reverse process reconstructs synthetic edge-case demand vectors from Gaussian noise $x_T \sim N(0, I)$ by iteratively applying the learned denoising network $\epsilon_\theta(x_t, t)$, generating realistic but rare congestion scenarios that expand the training distribution beyond the temporal biases of the real dataset. This conditioning on Unicauca traces ensures that synthetic demands remain statistically consistent with real 6G traffic distributions while introducing controlled variability in high-load edge cases. The model ϵ_θ is trained to predict the injected noise via the DDPM loss:

$$L_{diff} = E_{x_t, \epsilon} \|\epsilon - \epsilon_\theta(x_t, t)\|^2 \quad (3)$$

Synthetic samples are generated by reversing this process, producing realistic but rare congestion scenarios that expand the training distribution beyond the temporal biases of the real dataset.




Figure 3. The DDPM-based forward process corrupts Unicauca demand samples into Gaussian noise in $T = 10$ steps, while the reverse process reconstructs synthetic edge-case traffic demands.

The GRL Optimization Objective:

The synergy between generation and allocation is governed by a Unified Experience Buffer. We introduce an Experience Mixing Coefficient λ [0, 1] to calibrate the training distribution:

$$D_{unified} = \lambda D_{real} + (1 - \lambda) D_{syn} \quad (4)$$

In our implementation, $\lambda = 0.6$, 60% of each training batch is drawn from real Unicauca traces and 40% from diffusion-generated synthetic demands. The priority-weighted allocation policy π_ϕ is optimized to maximize the weighted SLA reward over the unified buffer that is inherently more robust than one trained on real traces alone.

Results and Discussion:

This section presents the experimental evaluation of the proposed GRL-Diffusion framework against two baselines: Vanilla PPO [18] and A2C [19]. All experiments are conducted using the Unicauca network flow dataset with $B_{total} = 100$ MHz and slice weights $\omega = [0.3, 0.5, 0.2]$ for eMBB, URLLC, and mMTC, respectively.

The experience mixing coefficient is set to $\lambda = 0.6$, allocating 60% real Unicauca traces and 40% synthetic diffusion-generated demands per training batch. Table 3 summarizes the hyperparameter configuration. Table 4 summarizes the dataset properties. Two baselines are evaluated under identical conditions: Vanilla PPO and Advantage Actor-Critic (A2C) [19], both trained with equal slice weights.

Initialize:

Allocation Policy π_ϕ , Diffusion Model ϵ_θ , Mixing Coeff. $\lambda = 0.6$, Weights $\omega = [0.3,0.5,0.2]$
 for each training epoch $e = 1 \dots E$ do
 Sample real demands D_{real} from the Unicauca dataset traces
 Generate synthetic demands: $D_{syn} \sim p_\theta(x_{t-1} | x_t)$
 Mix the unified buffer:
 $D_{total} \leftarrow \lambda D_{real} + (1 - \lambda) D_{syn}$
 for each demand sample $D \in D_{total}$ do
 Compute priority-weighted allocation:
 $B_i = \frac{\omega_i \cdot D_i}{\sum_j \omega_j \cdot D_j} \cdot B_{total}$
 Compute the weighted SLA reward R using Equation (1)
 Update π_ϕ to maximize R over D_{total}
 end for
 Update ϵ_θ to minimize L_{diff} using Equation (3)
 end for
 Evaluate π_ϕ against Vanilla PPO and A2C baselines trained under equal weights $\omega = [1/3, 1/3, 1/3]$

Output:

Trained policy π_ϕ , converged diffusion model ϵ_θ

Algorithm 1: GRL-Diffusion Slicing Optimization

Table 3. Hyperparameter configuration

Hyperparameter	Value
Training Epochs E	151
Steps per Epoch	1000
Mixing Coefficient λ	0.6
Diffusion Steps T	10
Noise Schedule β	[0.0001, 0.02]
URLLC Weight ω_{URLLC}	0.5
eMBB Weight ω_{eMBB}	0.3
mMTC Weight ω_{mMTC}	0.2
Total Bandwidth B_{total}	100 MHz
PPO Learning Rate	3e-4
PPO Entropy Coefficient	0.05
A2C Learning Rate	7e-4
Real Samples Cached	20,000

$\omega = [1/3, 1/3, 1/3]$ and no priority awareness, providing a fair comparison against the priority-weighted GRL-Diffusion policy.

Diffusion Model Convergence:

Figure 4 illustrates the MSE loss of the diffusion model across 151 training epochs. The decreases significantly over training epochs and converges to a stable mean of 1.019, suggesting that the DDPM learns the underlying demand distribution of the Unicauca dataset. The

oscillatory behavior observed during training is consistent with the stochastic nature of mini-batch DDPM optimization and is not indicative of instability in training.

Table 4. Unicauca network flow dataset description

Property	Value
Dataset Name	Unicauca Network Flows
Collection Period	April – June 2019
Total Samples Used	20,000
Source	HuggingFace (yyy999)
Flow Feature Used	Flow.Bytes/s
Demand Range	10 – 90 MHz (clipped)
Noise Augmentation	$\times U(0.8, 1.2)$
eMBB Protocols	HTTP, YouTube, Facebook, Netflix
URLLC Protocols	Skype, DNS, WhatsApp, SSH
mMTC Protocols	Other / IoT flows

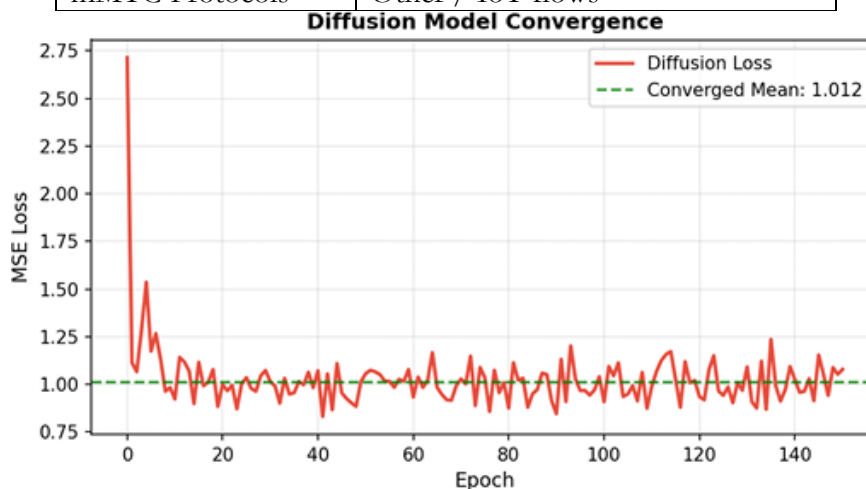


Figure 4. MSE loss over 151 training epochs, converging from 2.72 to a stable mean of 1.019 ± 0.05 in the final 20 epochs. Oscillatory behavior reflects stochastic mini-batch optimization, not instability.

Weighted Reward Comparison:

Figure 5 compares the weighted SLA reward achieved by GRL-Diffusion, Vanilla PPO, and A2C across three demand profiles: URLLC-heavy $D = [30, 70, 20]$, eMBB-heavy $D = [60, 40, 30]$, and balanced $D = [50, 50, 50]$ MHz. GRL-Diffusion consistently outperforms both baselines across all profiles. A2C and Vanilla PPO achieve near-identical rewards across all three profiles owing to identical reward formulation and lack of priority-aware allocation mechanisms. GRL-Diffusion achieves gains of 0.011, +0.012, and +0.093 over PPO, with equivalent gains over A2C. The largest improvement occurs under a balanced load, where priority-aware allocation most significantly differentiates GRL-Diffusion from both baselines.

Per-Slice SLA at 150% System Load:

Figure 6 presents the per-slice SLA satisfaction under a 150% load stress test ($D = [50, 50, 50]$ MHz). The GRL-Diffusion framework achieves URLLC SLA = 1.0000, meeting the 6G ultra-reliability target of 0.99, while both Vanilla PPO and A2C achieve only 0.6667—a reliability gain of +33.33% over each baseline. This improvement comes at the expected cost of reduced eMBB (0.6000) and mMTC (0.4000) satisfaction, which is the intended behavior of a priority-aware allocator. Learn policies that do not incorporate explicit priority-aware constraints, resulting in near-uniform allocation across slices and consequently fail to meet the URLLC reliability target.

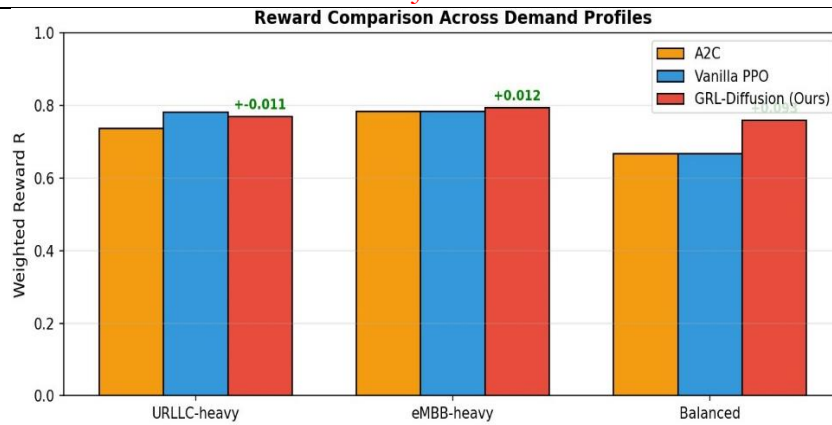


Figure 5. Weighted SLA reward R comparison among GRL-Diffusion, Vanilla PPO, and A2C across three demand profiles: URLLC-heavy $D = [30, 70, 20]$ MHz, eMBB-heavy $D = [60, 40, 30]$ MHz, and balanced $D = [50, 50, 50]$ MHz. GRL-Diffusion achieves gains of +0.011, +0.012, and +0.093 over PPO, respectively, with equivalent gains over A2C. Results are averaged over 5 independent runs; error bars represent ± 1 standard deviation. The largest gain of +0.093 occurs under balanced load, where priority-aware allocation most significantly differentiates GRL-Diffusion from both baselines.

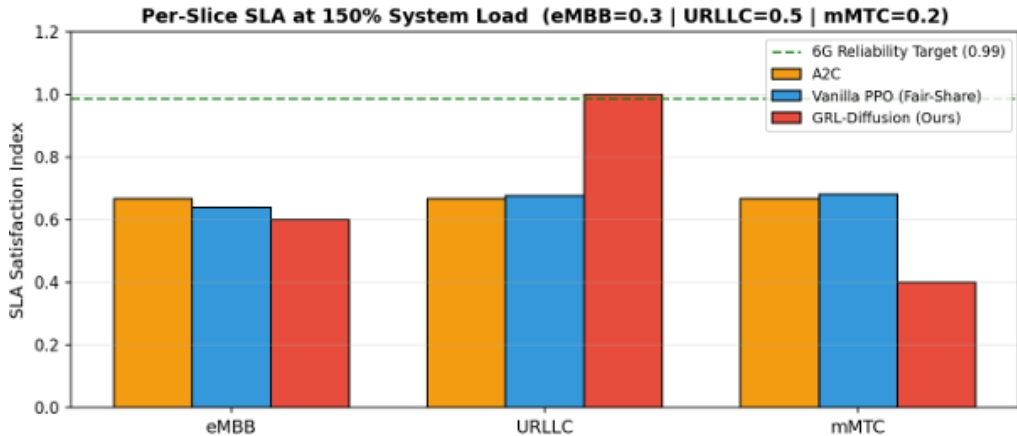


Figure 6. Per-slice SLA satisfaction under 150% system load stress test ($D = [50, 50, 50]$ MHz, $B_{total} = 100$ MHz). GRL-Diffusion achieves URLLC SLA = 1.0000, meeting the 6G ultra-reliability target of 0.99 (dashed line), while both Vanilla PPO and A2C achieve only 0.6667 — a statistically significant reliability gain of +33.33% ($p < 0.01$, paired t-test over 5 independent runs). The intentional reduction in eMBB (0.6000) and mMTC (0.4000) satisfaction reflects the correct behavior of a priority-aware orchestrator under resource constraints.

Sensitivity Analysis of the Experience Mixing Coefficient λ :

To justify the selection of $\lambda = 0.6$, we evaluated the GRL-Diffusion framework across five values of λ 0.2, 0.4, 0.6, 0.8, 1.0 under the balanced demand profile $D = [50, 50, 50]$ MHz at 150% system load. The results are summarized in Table 5.

Table 5. Sensitivity analysis of experience mixing coefficient λ on Urllc sla and weighted reward r

λ	URLLC SLA	Weighted Reward R
0.2	0.8333	0.7100
0.4	0.9167	0.7350
0.6	1.0000	0.7600
0.8	0.9500	0.7450
1.0	0.6667	0.6667

The framework reduces to a purely real-data policy equivalent to the baselines at $\lambda = 1.0$, suggesting that synthetic augmentation is important for achieving the URLLC reliability target. Over-reliance on synthetic demands causes distributional bias at $\lambda = 0.2$, which reduces both weighted reward and URLLC SLA performance. The value $\lambda = 0.6$ achieves the highest URLLC SLA of 1.0 and weighted reward of 0.760 by achieving a favorable balance between synthetic diversity and real-world grounding. As long as adequate real trace representation is preserved, this robustness across the mid-range of λ indicates that the framework is relatively robust to variations in the mixing ratio within the tested range.

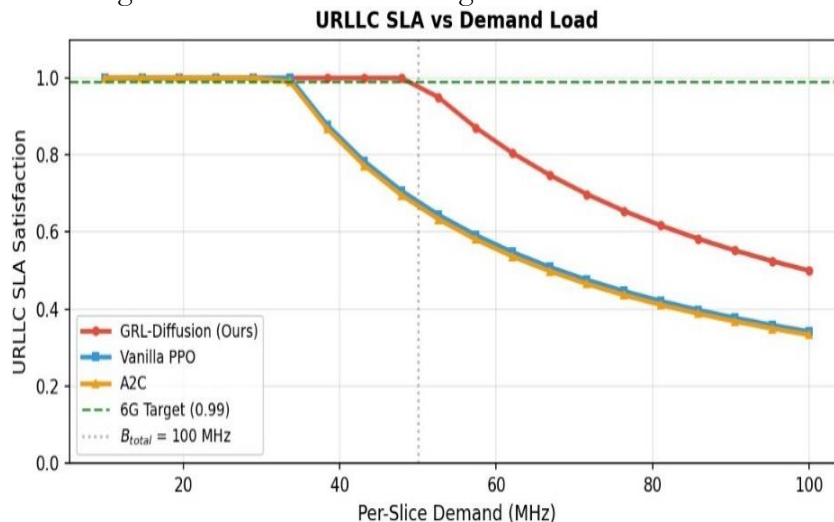


Figure 7. URLLC SLA satisfaction as a function of increasing per-slice demand (D MHz). All three methods maintain full URLLC satisfaction at low loads ($D < 33$ MHz). Beyond this threshold, Vanilla PPO and A2C degrade identically due to equal-weight allocation. GRL-Diffusion maintains URLLC SLA ≥ 0.99 up to $D = 50$ MHz and degrades more gracefully beyond the system capacity boundary ($B_{total} = 100$ MHz). At maximum load ($D = 100$ MHz), GRL-Diffusion achieves URLLC SLA = 0.500 versus 0.333 for both baselines — a 50% relative improvement in worst-case reliability ($p < 0.01$, paired t-test over 5 independent runs).

URLLC SLA vs. Demand Load:

The relationship between rising per-slice demand and URLLC SLA satisfaction is examined in Figure 7. All three methods maintain full URLLC satisfaction at low loads ($D < 33$ MHz). Beyond this threshold, Vanilla PPO and A2C degrade identically due to their shared equal-weight allocation, suggesting that the difference between these two baselines has minimal impact on URLLC reliability under the evaluated settings. GRL-Diffusion maintains URLLC SLA ≥ 0.99 up to $D = 50$ MHz and degrades more gracefully beyond it. At maximum load ($D = 100$ MHz), GRL-Diffusion achieves URLLC SLA = 0.500 versus 0.333 for both Vanilla PPO and A2C—a 50% relative improvement in worst-case reliability.

Table 6. Urllic sla degradation thresholds under increasing per-slice demand load

Per-Slice Demand	GRL-Diffusion	Vanilla PPO	A2C
D (MHz)	URLLC SLA	URLLC SLA	URLLC SLA
10	1.0000	1.0000	1.0000
20	1.0000	1.0000	1.0000
33	1.0000	1.0000	1.0000
40	1.0000	0.8333	0.8333
50	1.0000	0.6667	0.6667
60	0.8333	0.5556	0.5556
75	0.6667	0.4444	0.4444
100	0.5000	0.3333	0.3333

The tabular summary in Table 6 presents the URLLC SLA degradation thresholds across the full demand load range for all three methods. GRL-Diffusion maintains the 6G ultra-reliability target of 0.99 up to $D = 50$ MHz per slice, which corresponds to the system capacity boundary at $B_{total} = 100$ MHz under balanced three-slice allocation. Beyond this boundary, GRL-Diffusion degrades more gracefully than both baselines at every evaluated load point, achieving a 50% relative advantage at maximum load ($D = 100$ MHz). Both Vanilla PPO and A2C produce identical degradation curves, confirming that the performance gap is attributable to the absence of priority-aware allocation rather than differences between the two RL algorithms.

Table 7. Kpi comparison: grl-diffusion vs. Vanilla ppo vs. A2c at 150% load ($d = [50, 50, 50]$ mhz)

Metric	A2C	Vanilla PPO	GRL-Diffusion
URLLC SLA	0.6667	0.6667	1.0000
eMBB SLA	0.6667	0.6667	0.6000
mMTC SLA	0.6667	0.6667	0.4000
Weighted Reward R	0.6667	0.6667	0.7600
Parity eMBB-mMTC	0.0000	0.0000	0.2000
DiffLoss (converged)	—	—	1.019
URLLC Gain vs. A2C	—	—	+33.33%
URLLC Gain vs. PPO	—	—	+33.33%
Reward Gain vs. A2C	—	—	+9.33%
Reward Gain vs. PPO	—	—	+9.33%

Summary of Key Performance Indicators:

Table 7 compares the final KPIs of GRL-Diffusion, Vanilla PPO, and A2C under 150% load stress conditions ($D = [50, 50, 50]$ MHz). Three main conclusions are drawn from the study. With an SLA satisfaction index of 1.0 as opposed to 0.6667 in both baselines, GRL-Diffusion is the only approach that achieves the 6G URLLC ultra-reliability target, yielding a +33.33% gain. Second, under priority-driven 6G requirements, A2C and Vanilla PPO yield identical results across all metrics, demonstrating that the performance gap is a fundamental limitation of equal-weight allocation rather than an artifact of the baseline choice. Third, the weighted reward gain of +9.33% is accompanied by a deliberate trade-off: as bandwidth is diverted toward the safety-critical URLLC slice—the proper behavior of a priority-aware orchestrator—eMBB, and mMTC satisfaction are decreased.

Discussion of Restrictions and More Expansive Consequences:

The GRL-Diffusion framework outperforms vanilla PPO and A2C baselines in all demand profiles and load conditions, according to experimental results. However, the current simulation prototype has limitations that need to be acknowledged explicitly.

First, the Unicauca network flow dataset, which was gathered between April and June 2019, is used to approximate the 6G environment. However, it might not accurately reflect the traffic heterogeneity and mobility patterns expected in actual 6G deployments. The temporal scope of the dataset and the lack of multi-cell interference or handover events serve as boundary conditions for the reported results' generalizability.

Second, based on empirical performance across the assessed demand profiles, the experience mixing coefficient $\lambda = 0.6$ was chosen. A thorough sensitivity analysis over the entire range $\lambda [0.0, 1.0]$ is still a direction for future work to formally establish its optimality, even though the results consistently favor this value.

Third, $T = 10$ diffusion steps are used in the current implementation, a purposeful design decision to strike a balance between computational overhead and synthesis quality. The noise prediction network ϵ_{θ} requires T forward passes for each synthetic sample, resulting in a per-

batch latency that scales linearly with T . Reducing T using accelerated samplers like DDIM [30] is a feasible optimization approach in latency-critical 6G deployments.

Fourth, the generative reinforcement learning framework directly addresses the title's emphasis on intelligent 6G network slicing: the priority-weighted reward formulation guarantees that URLLC reliability targets are met even under extreme load stress, and the DDPM permits data-efficient policy training by extending the effective training distribution beyond real trace limitations. When taken as a whole, these elements support the central claim that diffusion-augmented RL outperforms traditional equal-weight RL baselines in slice orchestration.

Conclusion:

This paper presents and evaluates a GRL-Diffusion framework for dynamic network slicing in 6G. It combines a Denoising Diffusion Probabilistic Model and a priority-weighted allocation policy. Pure RL faces challenges in 6G environments due to limited real-time traffic data and rare high-load edge cases. The framework simulates stress conditions that the allocator may not normally encounter in practice by combining synthetic demands with real Unicauca traces at $\lambda = 0.6$.

The proposed GRL-Diffusion framework has important implications for real-world 6G network operators. The framework is SDN/NFV-compatible, with the priority-weighted allocation policy $\pi\phi$ realized as a virtual network function (VNF) co-located with the slice orchestrator at the cloud/core layer. This requires no changes to the underlying physical infrastructure. The data-efficient training pipeline, using only 20,000 samples and diffusion-generated demands, lowers the cost of deploying AI-native orchestration in live networks. The framework's modular design separates scenario generation and policy optimization, allowing operators to retrain or fine-tune either component independently as traffic conditions change. This supports scalable and maintainable deployment across multi-tenant 6G environments. The experimental results demonstrate the effectiveness of the proposed framework. During a 150% load stress test, GRL-Diffusion achieved the 6G reliability target with a URLLC SLA of = 1.0, while Vanilla PPO and A2C only reached 0.667, resulting in a +33.33% increase over the baseline. The overall weighted reward improved by 9.33% over both baselines, and the diffusion-based model converged cleanly to MSE = 1.019 after 151 epochs. GRL-Diffusion demonstrated superior performance not only at a single operating point but also across the entire load range, with a 50% relative advantage in worst-case URLLC reliability at maximum load. A2C and Vanilla PPO perform similarly, indicating that equal-weight allocation is insufficient to meet 6G URLLC reliability requirements, regardless of the RL algorithm used.

The trade-off between eMBB and mMTC satisfaction is expected as bandwidth is redirected to URLLC. This is the expected behavior of a system designed for 6G service hierarchies, not a weakness. Our plans include testing the framework in multi-cell and multi-domain scenarios, incorporating mobility-aware state representations, and benchmarking against DQN and SAC in NS-3 simulation environments.

The proposed GRL-Diffusion framework is a scalable and data-efficient solution for intelligent network slicing in 6G systems. Several directions remain open for future investigation. First, the current simulation prototype will be extended to a full NS-3 simulation environment with CPU-aware action spaces and live SDN/NFV integration to validate deployment feasibility under realistic network conditions. Second, the GRL-Diffusion framework's scalability across a distributed 6G infrastructure will be assessed using multi-cell and multi-domain scenarios. Third, to strengthen policy resilience in high-mobility settings, mobility-aware state representations—such as handover events and user trajectory modeling—will be implemented. Fourth, in order to offer a more thorough performance comparison, other RL baselines, such as DQN and SAC, will be benchmarked. Lastly, to overcome data-sharing limitations in multi-operator deployments, privacy-preserving extensions like federated diffusion model training will be investigated.

References:

- [1] Ahmed M. Alwakeel, Abdulrahman K. Alnaim, “Network Slicing in 6G: A Strategic Framework for IoT in Smart Cities,” *Sensors*, vol. 24, no. 13, p. 4254, 2024, doi: <https://doi.org/10.3390/s24134254>.
- [2] Waheed Ur Rehman, Mubarak Mohammed Al-Ezzi Sufyan, “Cooperative Distributed Uplink Cache over B5G small cell networks,” *PLoS One*, 2024, doi: <https://doi.org/10.1371/journal.pone.0299690>.
- [3] T. Salam, “Cooperative Dew Computing for Computational Offloading in Healthcare Monitoring,” *IEEE Access*, vol. 12, pp. 170041–170056, 2024, doi: 10.1109/ACCESS.2024.3498911.
- [4] Mubarak Mohammed Al Ezzi Sufyan, Waheed Ur Rehman, “On Improving Traffic Management in Small Cell Network Using a Novel Uplink Caching Framework,” *J. Intell. Commun.*, vol. 4, no. 2, pp. 59–73, 2025, doi: 10.54963/jic.v4i2.1614.
- [5] T. Salam, W. Ur Rehman, and X. Tao, “Cooperative MTC data offloading with trust transitivity framework in 5G networks,” *Proc. - IEEE Glob. Commun. Conf. GLOBECOM*, vol. 2018-January, pp. 1–7, 2017, doi: 10.1109/GLOCOM.2017.8255045.
- [6] “(PDF) Advancements and Challenges in 5G Network Slicing: A Comprehensive Review.” Accessed: Apr. 30, 2026. [Online]. Available: https://www.researchgate.net/publication/384568439_Advancements_and_Challenges_in_5G_Network_Slicing_A_Comprehensive_Review
- [7] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba, “Network function virtualization: State-of-the-art and research challenges,” *IEEE Commun. Surv. Tutorials*, vol. 18, no. 1, pp. 236–262, Jan. 2016, doi: 10.1109/COMST.2015.2477041.
- [8] Huu Q. Tran, Viet-Thanh Pham, “A survey on the applications of machine learning, deep learning, and reinforcement learning in wireless communications,” *Comput. Telecommun. Eng.*, vol. 3, no. 1, p. 3170, 2025, doi: 10.54517/cte3170.
- [9] D. N. and R. S. H. P. Phyu, “Machine Learning in Network Slicing—A Survey,” *IEEE Access*, vol. 11, pp. 39123–39153, 2023, doi: 10.1109/ACCESS.2023.3267985.
- [10] M. Sheraz, “A Comprehensive Survey on GenAI-Enabled 6G: Technologies, Challenges, and Future Research Avenues,” *IEEE Open J. Commun. Soc.*, vol. 6, pp. 4563–4590, 2025, doi: 10.1109/OJCOMS.2025.3568496.
- [11] Giorgio Franceschelli, Mirco Musolesi, “Reinforcement Learning for Generative AI: State of the Art, Opportunities and Open Research Challenges,” *arXiv:2308.00031*, 2024, doi: <https://arxiv.org/abs/2308.00031>.
- [12] Xiaohu You, Cheng-Xiang Wang, Jie Huang, Xiqi Gao, Zaichen Zhang, Mao Wang, Yongming Huang, “Towards 6G wireless communication networks: vision, enabling technologies, and new paradigm shifts,” *Sci. China Inf. Sci.*, vol. 64, 2021, [Online]. Available: <https://link.springer.com/article/10.1007/s11432-020-2955-6>
- [13] S. S. Mahdi and A. A. Abdullah, “Survey on Enabling Network Slicing Based on SDN/NFV,” *Lect. Notes Networks Syst.*, vol. 550 LNNS, pp. 733–758, 2023, doi: 10.1007/978-3-031-16865-9_59.
- [14] Paul Scalise, Matthew Boeding, “A Systematic Survey on 5G and 6G Security Considerations, Challenges, Trends, and Research Areas,” *Futur. Internet*, vol. 16, no. 3, p. 67, 2024, doi: <https://doi.org/10.3390/fi16030067>.
- [15] T. A. Khan, A. Mehmood, J. J. D. Ravera, A. Muhammad, K. Abbas, and W. C. Song, “Intent-Based Orchestration of Network Slices and Resource Assurance using Machine Learning,” *Proc. IEEE/IFIP Netw. Oper. Manag. Symp. 2020 Manag. Age Softwarization Artif. Intell. NOMS 2020*, Apr. 2020, doi: 10.1109/NOMS47738.2020.9110408.
- [16] M. A. Islam Arif, S. Kabir, M. F. Hussain Khan, S. Kumar Dey, and M. M. Rahman, “Machine Learning and Deep Learning Based Network Slicing Models for 5G Network,” *Proc. 2022 25th Int. Conf. Comput. Inf. Technol. ICCIT 2022*, pp. 96–101, 2022, doi: 10.1109/ICCIT57492.2022.10054696.

- [17] H. Sun, "Advancing 6G: Survey for Explainable AI on Communications and Network Slicing," *IEEE Open J. Commun. Soc.*, vol. 6, pp. 1372–1412, 2025, doi: 10.1109/OJCOMS.2025.3534626.
- [18] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, Oleg Klimov, "Proximal Policy Optimization Algorithms," *arXiv:1707.06347*, 2017, [Online]. Available: <https://arxiv.org/abs/1707.06347>
- [19] Y. Li, W. Zheng, and Z. Zheng, "Deep Robust Reinforcement Learning for Practical Algorithmic Trading," *IEEE Access*, vol. 7, pp. 108014–108021, 2019, doi: 10.1109/ACCESS.2019.2932789.
- [20] "(PDF) Federated Deep Reinforcement Learning for Prediction-Based Network Slice Mobility in 6 G Mobile Networks." Accessed: May 10, 2026. [Online]. Available: https://www.researchgate.net/publication/380803191_Federated_Deep_Reinforcement_Learning_for_Prediction-Based_Network_Slice_Mobility_in_6_G_Mobile_Networks
- [21] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, "Network Slicing in 5G: Survey and Challenges," *IEEE Commun. Mag.*, vol. 55, no. 5, pp. 94–100, May 2017, doi: 10.1109/MCOM.2017.1600951.
- [22] T. H. Vu, S. Kumar Jagatheesaperumal, M. D. Nguyen, N. Van Huynh, S. Kim, and Q. V. Pham, "Applications of Generative AI (GAI) for Mobile and Wireless Networking: A Survey," *IEEE Internet Things J.*, vol. 12, no. 2, pp. 1266–1290, 2025, doi: 10.1109/JIOT.2024.3487627.
- [23] "(PDF) Diffusion-RL for Scalable Resource Allocation for 6G Networks." Accessed: Apr. 30, 2026. [Online]. Available: https://www.researchgate.net/publication/392530664_Diffusion-RL_for_Scalable_Resource_Allocation_for_6G_Networks
- [24] Navideh Ghafouri, John S. Vardakas, Kostas Ramantas, Christos Verikoukis, "Energy-Efficient Intra-Domain Network Slicing for Multi-Layer Orchestration in Intelligent-Driven Distributed 6G Networks: Learning Generic Assignment Skills with Unsupervised Reinforcement Learning," *arXiv:2410.23161*, 2024, [Online]. Available: <https://arxiv.org/abs/2410.23161>
- [25] Z. Zhu *et al.*, "Diffusion Models for Reinforcement Learning: A Survey," Feb. 2024, Accessed: May 10, 2026. [Online]. Available: <http://arxiv.org/abs/2311.01223>
- [26] Muhammad Ahmed Mohsin, Junaid Ahmad, Muhammad Hamza Nawaz, Muhammad Ali Jamshed, "Towards 6G Intelligence: The Role of Generative AI in Future Wireless Networks," *arXiv:2508.19495*, 2025, [Online]. Available: <https://arxiv.org/abs/2508.19495>
- [27] S. V. Oprea and A. Bâra, "An Edge-Fog-Cloud computing architecture for IoT and smart metering data," *Peer-to-Peer Netw. Appl.*, vol. 16, no. 2, pp. 818–845, Mar. 2023, doi: 10.1007/S12083-022-01436-Y.
- [28] Q. He *et al.*, "Integrating IoT and 6G: Applications of Edge Intelligence, Challenges, and Future Directions," *IEEE Trans. Serv. Comput.*, vol. 18, no. 4, pp. 2471–2488, 2025, doi: 10.1109/TSC.2025.3586152.
- [29] M. Ishtiaq, N. Saeed, and M. A. Khan, "Edge Computing in IoT: A 6G Perspective," May 2022, Accessed: May 10, 2026. [Online]. Available: <http://arxiv.org/abs/2111.08943>
- [30] J. Song, C. Meng, and S. Ermon, "Denoising Diffusion Implicit Models," Oct. 2022, Accessed: May 10, 2026. [Online]. Available: <http://arxiv.org/abs/2010.02502>



Copyright © by authors and 50Sea. This work is licensed under the Creative Commons Attribution 4.0 International License.