



Original Research Paper

DQN-Based Channel-Adaptive Federated Learning over Unreliable Wireless Channels with Hybrid Gradient Compression

M. Kaliappan¹, S Vimal², Gaurav Dhiman^{3,4,5}, M. Jayanthi⁶

¹Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India

²Department of Computer Science and Engineering, KIT- Kalaighnar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India

³Department of Computer Science and Engineering, Yuan Ze University, Tao Yuan, Taiwan

⁴University Centre for Research and Development, Chandigarh University, Mohali, India

⁵Division of Research and Development, Lovely Professional University, Phagwara, India

⁶Department of Cybersecurity, KIT- Kalaighnar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India

Email: kaliappan@ritripm.ac.in, svimal@kitcbe.ac.in, gdhiman0001@gmail.com, jayanthikit2020@gmail.com

ABSTRACT

Federated learning (FL) over wireless networks suffers from two compounding degradations: stochastic channel fading that induces gradient transmission failures, and fixed-rate compression policies that either waste bandwidth under good channel conditions or sacrifice model fidelity under poor ones. This paper proposes a Deep Q-Network (DQN)-based channel-adaptive FL framework that dynamically selects the optimal gradient compression policy jointly controlling top-k sparsification ratio and quantization bit-width—for each participating client in every communication round, conditioned on real-time channel state information (CSI). The DQN agent observes a three-dimensional state comprising instantaneous SNR, achievable data rate, and local data volume, and outputs one of four discrete compression actions ranging from full-precision 16-bit transmission to aggressive 4-bit quantization with 90% sparsification. A wireless channel simulator models Rayleigh fading, bandwidth-limited Shannon capacity, and outage events, providing a realistic training environment for the agent. Experiments on the Fashion-MNIST benchmark with non-IID client data partitioning demonstrate that the proposed method achieves 93.18% test accuracy a gain of 4.37–8.86 percentage points over three strong baselines—while reducing total transmission energy by 55.6% and communication overhead by 58.9% compared to FedAvg without compression. The DQN agent converges to a coherent channel-aware policy: selecting high-quality (low-compression) actions for clients with strong channel conditions and aggressive compression for outage-prone clients, thereby maximizing successful participation. These results establish the proposed framework as a principled, resource-efficient solution for deploying FL in practical wireless communication systems.

Keywords: *Federated Learning, Deep Q-Network, Channel-Adaptive Compression, Wireless Communications, Gradient Sparsification, Quantization, Non-IID, Fashion-MNIST*

*Corresponding Author:

Received: 25th July, 2026 | Revised: 26th August, 2026 | Accepted: 18th September, 2026 | Available Online: 24th September, 2026

(DOI): [10.70102/AEJ.2026.18.7s.47](https://doi.org/10.70102/AEJ.2026.18.7s.47)

I. INTRODUCTION

The proliferation of intelligent edge devices—smartphones, autonomous vehicles, industrial IoT sensors, and wearable health monitors—has generated vast quantities of sensitive, heterogeneous data at the network periphery. Federated learning (FL) [1] addresses the dual imperative of leveraging this data for machine learning model training while respecting data privacy constraints, by aggregating locally computed gradient updates rather than raw data. However, deploying FL in realistic wireless environments exposes a critical and underexplored gap between the idealized assumptions of FL theory and the harsh realities of wireless propagation.

Wireless channels exhibit time-varying fading, interference, and stochastic outages that fundamentally undermine the reliable gradient exchange upon which FL convergence depends. When a client experiences a deep fade event, its gradient update—computed through an entire epoch of local training—is either lost entirely or must be retransmitted at considerable energy cost. For battery-constrained edge devices operating in crowded spectrum environments, this represents an unacceptable waste of both computational and communicational resources. Simultaneously, transmitting high-dimensional neural network gradients over bandwidth-limited wireless links incurs prohibitive communication overhead, often constituting the dominant bottleneck in FL system throughput.

Three interrelated threats challenge the viability of wireless FL systems. First, channel stochasticity introduces non-deterministic participation patterns: clients with poor instantaneous channel conditions cannot reliably contribute in each round, creating systematic bias in the global model update. Second, one-size-fits-all compression policies—whether applied uniformly across all clients or fixed as a system-wide constant—fail to exploit the dynamic diversity of wireless channel conditions. A client experiencing SNR of 28 dB can afford to transmit high-fidelity gradients, while one at 3 dB SNR must either be excluded or severely compress its update. Third, the joint optimization of compression policy, client selection, and aggregation weighting under stochastic channel dynamics constitutes a non-convex, non-stationary optimization problem that is intractable by classical methods.

Deep reinforcement learning (DRL) offers a compelling paradigm for this optimization challenge. A DRL agent interacting with the wireless FL environment can learn, through trial and reward feedback, a channel-aware policy that maps observed system state to compression actions—without requiring explicit knowledge of the channel distribution, model architecture, or convergence rate. This paper exploits this opportunity by formulating the compression-selection

problem as a Markov Decision Process (MDP) and training a DQN agent to solve it within the FL training loop.

The scope of this work encompasses: (i) a complete channel-adaptive FL system design integrating wireless channel simulation, DQN-based policy learning, hybrid gradient compression, and weighted federated aggregation; (ii) comprehensive experimental evaluation on Fashion-MNIST with non-IID data across ten heterogeneous clients; (iii) analysis of learned DQN policies as a function of channel conditions; and (iv) comparison against three competitive baselines spanning the compression design space. The principal contributions of this paper are summarized as follows:

- A novel DQN-based adaptive compression framework for wireless FL that jointly optimizes gradient sparsification and quantization bit-width per client per communication round, conditioned on real-time CSI.
- A realistic wireless channel simulation model incorporating Rayleigh fading, Shannon capacity bounds, and probabilistic outage events, seamlessly integrated into the FL training loop.
- A carefully designed reward function that captures the multi-objective nature of wireless FL—balancing model accuracy improvement, energy efficiency, and communication overhead reduction.
- Extensive empirical evaluation demonstrating 93.18% test accuracy on Fashion-MNIST, with 55.6% energy reduction and 58.9% communication overhead reduction versus the FedAvg baseline.
- Visualization and interpretation of the emergent DQN compression policy, revealing its channel-aware characteristics and validating the theoretical design intuitions.

The remainder of this paper is organized as follows. Section II surveys related work from 2021 to 2025. Section III presents the system model and mathematical formulations. Section IV describes the proposed DQN-FL framework. Section V presents experimental results. Section VI compares with prior works. Section VII concludes the paper with directions for future research.

II. LITERATURE REVIEW

This section surveys the existing body of research spanning wireless federated learning, gradient compression, and reinforcement learning-based resource management. Table I provides a structured overview of the 35 most relevant references.

The foundational FedAvg algorithm [1] established the communication-efficient FL paradigm by replacing gradient transmission with full local model updates. The FedProx framework [23] extended this with a proximal regularization term to handle systems heterogeneity, proving convergence guarantees even with partial participation. SCAFFOLD [24] introduced variance reduction through control variates, addressing the client drift problem that plagues non-IID data distributions. While these works established theoretical FL foundations, they uniformly assumed reliable, high-bandwidth communication channels—an assumption that critically fails in wireless deployments. The non-IID convergence analysis of [9] provided tight bounds on accuracy degradation under data heterogeneity, motivating the channel-aware client selection strategy embedded in our DQN agent.

The seminal work on gradient sparsification [4] demonstrated that transmitting only the top-k magnitude gradients with error feedback can achieve near-lossless convergence with 99% communication reduction on fully connected networks. Quantization-aware FL [7] extended this to adaptive bit-width selection from 1 to 8 bits, demonstrating that low-bit quantization incurs minimal accuracy loss when combined with stochastic rounding. The survey [14] comprehensively taxonomized sparse communication methods for FL, identifying that hybrid approaches combining sparsification and quantization consistently outperform either technique in isolation—a finding directly exploited by our compression module. Gradient compression with error feedback [25] provided theoretical analysis showing that biased compressors can still converge provided the compression error is bounded, a property satisfied by our top-k sparsification with feedback accumulation.

Channel-aware gradient compression [5] was among the first works to condition the compression ratio on wireless channel state, demonstrating 25% communication reduction on Fashion-MNIST. However, this work used a fixed SNR threshold policy rather than a learned adaptive policy, limiting its generalization across channel distributions. Our DQN agent learns a continuous policy over the SNR space, exploiting nonlinear relationships between channel quality and achievable compression ratio.

DRL has emerged as the dominant paradigm for online resource management in wireless networks. The seminal application of DQN to heterogeneous network resource allocation [2] demonstrated that DRL agents can learn near-optimal scheduling policies in non-stationary environments without explicit system model knowledge. DQN-based power control [6] applied this paradigm to LTE HetNets, achieving significant energy savings versus model predictive control baselines. Double DQN [28] addressed the Q-value overestimation

bias inherent in standard DQN, improving stability in high-dimensional wireless resource optimization. Actor-critic approaches [12] extended DRL to continuous action spaces relevant for power control, while the multi-armed bandit approach [20] provided a low-complexity alternative suitable for compression policy selection with a small discrete action set—a setting well-matched to our four-action DQN.

The integration of DRL specifically within FL systems has accelerated recently. The energy-aware FL framework [21] combined DRL with edge caching to minimize round-trip latency and energy cost simultaneously. DRL-based joint user scheduling and beamforming [33] addressed the multi-antenna FL scenario where spatial multiplexing can serve multiple clients simultaneously. Attention-based client selection [32] applied transformer architectures to learn which clients to include in each round based on gradient diversity and channel quality—a complementary approach to our compression-level selection.

Energy efficiency has emerged as a primary design criterion for wireless FL deployments on battery-powered devices. The joint optimization framework [3] formulated energy-efficient FL as a mixed-integer nonlinear program, jointly optimizing device selection, local computation frequency, and uplink power allocation. The joint bandwidth and power optimization [15] extended this to include channel-dependent achievable rates, solving the resulting non-convex problem with successive convex approximation. Hierarchical FL [30] introduced an edge-cloud hierarchy that reduces long-range communication energy by performing intermediate aggregation at edge servers. The model compression survey [31] benchmarked energy implications of different compression techniques across edge hardware platforms, revealing that the energy savings from reduced communication dominate the overhead of compression computation for networks larger than 10K parameters—directly applicable to our SmallCNN architecture.

Over-the-air computation (AirComp) [11] exploits the superposition property of wireless channels to perform analog aggregation, avoiding digital transmission entirely. While theoretically elegant, AirComp requires tight synchronization and analog front-end hardware modifications impractical for commodity devices. Semantic communication [34] for FL updates represents an emerging paradigm that compresses model updates to task-relevant semantic features rather than numerical gradients, offering potential for extreme compression ratios. Robust FL [22] addressed adversarial channel conditions and Byzantine clients simultaneously, introducing a geometric median aggregation robust to malicious gradient injection. Outage-resilient federated distillation [18] used knowledge distillation rather than gradient aggregation

to tolerate high outage rates, though at the cost of requiring a public dataset at the server[36,37,38].

The convergence of DRL-based adaptive FL has been formally analyzed in [35], which proves that the DRL-FL system converges to a stationary point of the global objective under standard Lipschitz smoothness and bounded gradient variance assumptions. This theoretical backing provides the mathematical foundation for our empirical approach and validates the optimality of the DQN-learned policy in the limit of sufficient training[39,40].

Despite the rich body of related work, a clear gap exists: no prior work simultaneously (i) learns a per-client per-round compression policy conditioned on multi-dimensional CSI using deep RL[41], (ii) integrates hybrid sparsification-quantization compression with a shared action space, and (iii) validates on Fashion-MNIST with realistic wireless channel simulation including outage events. Our proposed DQN channel-adaptive FL framework fills this gap, as detailed in the following sections [42]-[63].

Table I: Summary of Related Literature (2017–2025)

R ef .	Title / Focus	Y ea r	Journal /Venue	Dataset	Key Contrib ution
[1]	FedAvg: Commu nication- efficient FL	20 17	AISTA TS	MNIST	Foundat ional federate d averagi ng
[2]	Deep reinforce ment learning in wireless network s	20 21	IEEE TWC	Syntheti c	DRL for resource allocati on
[3]	Energy- efficient FL over wireless	20 21	IEEE JSAC	CIFAR- 10	Joint device selectio n and

					compres sion
[4]	Commu nication compres sion for distribut ed learning	20 21	IEEE TSP	ImageN et	Top-k sparsific ation framew ork
[5]	Channel -aware gradient compres sion in FL	20 22	IEEE TWC	FMNIS T	SNR- depende nt compres sion
[6]	DQN- based resource manage ment in HetNets	20 22	IEEE Networ k	LTE traces	Q- learning for scheduli ng
[7]	Quantiza tion- aware FL under bandwid th constrai nts	20 22	IEEE TCOM	CIFAR- 10	1-bit to 8-bit adaptive quant.
[8]	Adaptiv e model aggregat ion with unreliabl e links	20 22	IEEE IoT-J	Fashion - MNIST	Asynchr onous FL for IoT

[9]	Wireless FL with non-IID data: convergence analysis	2022	IEEE TNNLS	MNIST/CIFAR	Non-IID convergence bounds
[10]	Gradient sparsification for bandwidth-limited FL	2022	IEEE TMC	EMNIST	Error-feedback sparsification
[11]	Over-the-air computation for FL	2023	IEEE TWC	Synthetic	AirComp aggregation method
[12]	Actor-critic RL for FL scheduling	2023	IEEE TMC	CIFAR-100	A3C-based client selection
[13]	Personalized FL with heterogeneous data	2023	ICLR	Multiple	Per-FedAvg algorithm
[14]	Sparse communication in FL: survey	2023	IEEE Comm. Surv.	Multiple	Comprehensive survey
[15]	Joint bandwidth	2023	IEEE TWC	Fashion MNIST	MINLP optimiz

	th and power optimization in wireless FL				ation framework
[16]	Model pruning for FL over heterogeneous devices	2023	IEEE TCOM	MNIST	Layer-wise pruning strategy
[17]	DRL-based power control in millimeter wave FL	2023	IEEE Access	Custom	mmWave channel-adaptive FL
[18]	Outage-resilient federated distillation	2023	IEEE GLOBE COM	CIFAR-10	Knowledge distillation under outage
[19]	Hybrid compression for FL: sparsity + quantization	2024	IEEE TCOM	ImageNet	Combined sparsification-quantization
[20]	Multi-armed	2024	IEEE TWC	FMNIST	Bandit-based

	bandit for FL compression selection				adaptive compression
[21]	Energy-aware FL with DRL and edge caching	2024	IEEE IoT-J	Synthetic	DRL + edge caching synergy
[22]	Robust FL under adversarial channel conditions	2024	IEEE TIFS	MNIST	Byzantine- robust aggregation
[23]	FedProx : proximal FL for heterogeneous systems	2020	MLSys	Multiple	FedProx convergence guarantees
[24]	Scaffold : stochastic controlled averaging for FL	2020	ICML	Multiple	Variance reduction in FL
[25]	Gradient compression	2021	NeurIPS	ImageNet	Biased compression

	with error feedback				sion analysis
[26]	Adaptive quantization via Bayesian optimization	2024	IEEE TNNLS	CIFAR-100	Bayesian bit-width selection
[27]	FL over heterogeneous wireless: a systems view	2024	IEEE Comm. Mag.	Multiple	System-level FL challenges
[28]	Double DQN for network resource optimization	2024	IEEE TMC	Real traces	DDQN stability improvements
[29]	Channel estimation for FL using pilot signals	2024	IEEE TWC	Synthetic	Pilot-based CSI for FL
[30]	Hierarchical FL with channel-aware aggregation	2024	IEEE JSAC	CIFAR-10	Edge-cloud FL hierarchy

[3 1]	Model compression survey for edge intelligence	20 24	IEEE Comm. Surv.	Multipl e	Survey of edge ML compression
[3 2]	Attention-based client selection in wireless FL	20 25	IEEE TWC	FMNIS T	Transformer client selection
[3 3]	DRL for joint user scheduling and beamforming in FL	20 25	IEEE TSP	Synthetic	MIMO FL with RL scheduling
[3 4]	Semantic communication for FL model updates	20 25	IEEE TCOM	Custom NLP	Semantic compression paradigm
[3 5]	Convergence of DRL-based adaptive FL	20 25	IEEE TNNLS	Multipl e	Theoretical convergence proof

Consider a FL system comprising a central server and $N = 10$ clients, indexed by $i \in \{1, 2, \dots, N\}$.

Each client i holds a local dataset D_i of $n_i = |D_i|$ labeled samples drawn from a heterogeneous, non-IID distribution P_i . The global objective is to minimize the weighted empirical loss:

$$F(w) = \sum_i (n_i / n) F_i(w), \quad \text{where } F_i(w) = (1/n_i) \sum_{(x,y) \in D_i} \ell(w; x, y) \quad (1)$$

where $w \in \mathbb{R}^d$ denotes the global model parameters, $n = \sum_i n_i$ is the total sample count, and $\ell(\cdot)$ is the cross-entropy loss. In each communication round t , the server selects a subset $S_t \subseteq \{1, \dots, N\}$ of $|S_t| = K = 5$ clients. Each selected client $i \in S_t$ receives the current global model w_t , performs $E = 1$ epoch of local SGD, and computes a gradient update $\Delta_{i,t} = w_{i,t}^{\text{local}} - w_t$.

The uplink wireless channel from client i to the server in round t is modeled as a Rayleigh flat-fading channel with path loss. The instantaneous received signal-to-noise ratio (SNR) is drawn as:

$$\gamma_i \sim \text{Uniform}(0, 30) \text{ dB}, \quad \gamma_{i,\text{linear}} = 10^{\gamma_i/10} \quad (2)$$

The achievable uplink data rate for client i is given by the Shannon capacity formula:

$$R_i = B \cdot \log_2(1 + \gamma_{i,\text{linear}}) \quad [\text{bits/s}] \quad (3)$$

where $B = 10^6$ Hz is the system bandwidth. A channel outage event for client i occurs with probability: $P_{\text{out},i} = \exp(-\gamma_i/8)$ (4)

This outage model captures the physical reality that low-SNR links have exponentially higher failure probability. A Bernoulli draw with parameter $1 - P_{\text{out},i}$ determines whether client i 's transmission succeeds in round t . Failed transmissions result in the client's update being discarded, and the client does not contribute to the aggregation in that round.

The gradient update $\Delta_{i,t} \in \mathbb{R}^d$ is compressed before wireless transmission. The compression pipeline consists of two sequential operations:

Given sparsification ratio $\rho \in (0, 1]$, the top- k operator retains only the $k = \lceil \rho \cdot d \rceil$ largest magnitude gradient components:

$$\text{Sparsify}(\Delta, \rho) = \arg \max_{\{S \subseteq [d], |S|=k\}} \sum_{j \in S} |\Delta_j| \quad (5)$$

The remaining $(d - k)$ components are set to zero, yielding a sparse tensor transmitted as a compressed sparse row (CSR) format.

The sparse gradient is subsequently quantized to b -bit precision. Let $\Delta_{\min}$ and $\Delta_{\max}$ denote the minimum and maximum values of the sparsified gradient. The quantization operation is:

$$\Delta_q = \text{round}((\Delta - \Delta_{\min}) / ((\Delta_{\max} - \Delta_{\min}) / (2^b - 1))) \cdot ((\Delta_{\max} - \Delta_{\min}) / (2^b - 1)) + \Delta_{\min} \quad (6)$$

The total communication cost in bits for client i with sparsity ρ and bit-width b is:

$$C_i(\rho, b) = \rho \cdot d \cdot b \quad [\text{bits}] \quad (7)$$

The transmission energy for client i to deliver its compressed update over the wireless channel is:

$$E_i = P_{i_tx} \cdot \tau_i = (P_{base} + \alpha \cdot \gamma_i) \cdot C_i(\rho, b) / R_i \quad [\text{Joules}] \quad (8)$$

where $P_{base} = 0.1$ W is the baseline transmit power, $\alpha = 0.01$ is a SNR-dependent power scaling coefficient, and $\tau_i = C_i/R_i$ is the transmission time. This model captures the empirical observation that devices operating in poor channel conditions must increase transmit power to maintain acceptable packet error rates.

The server aggregates the compressed updates from successful clients $S_{t_success} \subseteq S_t$ using weighted FedAvg:

$$w_{\{t+1\}} = w_t + \sum_{i \in S_{t_success}} (n_i / \sum_{j \in S_{t_success}} n_j) \cdot \tilde{\Delta}_i \quad (9)$$

where $\tilde{\Delta}_i$ denotes the decompressed gradient update received from client i . The compression-decompression pipeline introduces a bounded error, analyzed in the convergence bound of [35].

3. PROPOSED DQN CHANNEL-ADAPTIVE FL FRAMEWORK

3.1. System Architecture Overview

The proposed framework integrates three co-designed components: (1) a wireless channel estimator that characterizes per-client channel conditions in each round, (2) a DQN-based compression policy agent that maps channel state to compression actions, and (3) a hybrid gradient compression module that implements the selected compression policy. Figure 1 depicts the complete system block diagram. The DQN agent receives multi-dimensional channel state information from the channel estimator and outputs per-client compression actions, which are executed by the adaptive compression module before wireless transmission to the global FL server.

This work introduces an innovative channel-adaptive federated learning framework that directly tackles a significant deficiency in current FL research: the implausible presumption of dependable wireless communication. The proposed system uses a DQN-based decision engine and real-time channel estimation (SNR, rate, outage) to choose both client participation and compression levels on the fly. This makes it suitable to use for a broad spectrum of wireless environments which can be unreliable, including varying signal-to-noise ratio and varying availability of the network. The proposal of an optimization scheme that takes into account all three factors mentioned, i.e., model accuracy, communication cost, and energy cost, is a remarkable improvement from the conventional FL approach. Additionally, implementing adaptive sparsification and quantization on the client side, tailored to specific network conditions, provides a practical and scalable approach for deploying edge computing in 5G/6G environments. This cross-layer design that connects machine learning, wireless communication, and resource

optimization is very new, useful, and in line with new research trends.

3.2. Markov Decision Process Formulation

The compression policy selection problem is formulated as a finite-horizon MDP with the following components:

State Space

The state vector for client i in round t is:

$$s_{i_t} = [\gamma_{i_t} / \gamma_{max}, R_{i_t} / R_{max}, n_i / n] \in [0, 1]^3 \quad (10)$$

This three-dimensional normalized state captures: (i) the instantaneous channel quality relative to the maximum observed SNR (30 dB), (ii) the achievable rate normalized by the peak rate across all clients in the current round, and (iii) the local data volume fraction, which encodes the relative importance of client i 's gradient update. Normalizing all state components to $[0, 1]$ ensures numerically stable DQN training.

Action Space

The DQN agent selects from a discrete set of $A = 4$ compression actions:

$$\mathcal{A} = \{a_0: \text{Skip}, a_1: (\rho=1.0, b=16), a_2: (\rho=0.3, b=8), a_3: (\rho=0.1, b=4)\} \quad (11)$$

Action a_0 (Skip) excludes a client from the current round, appropriate for extremely poor channel conditions where any transmission would be unreliable. Action a_1 transmits the full gradient in 16-bit precision, suitable for excellent channel conditions. Actions a_2 and a_3 apply progressively more aggressive compression, trading gradient fidelity for communication efficiency.

Reward Function

The scalar reward signal in round t is designed to balance three competing objectives:

$$r_t = \lambda_1 \cdot \Delta Acc_t - \lambda_2 \cdot E_{t_total} - \lambda_3 \cdot C_{t_total} \quad (12)$$

where $\Delta Acc_t = Acc_t - Acc_{\{t-1\}}$ is the incremental test accuracy gain, $E_{t_total} = \sum_i E_{i_t}$ is the total transmission energy in round t , and $C_{t_total} = \sum_i C_i(\rho_i, b_i)$ is the total communication overhead. The hyperparameters $\lambda_1 = 20.0$, $\lambda_2 = 10^{-6}$, and $\lambda_3 = 10^{-9}$ are set such that a 1% accuracy gain is weighted equivalently to approximately 20 J of energy and 20 Gbits of overhead, prioritizing model quality while penalizing resource waste.

3.3. DQN Agent Architecture and Training

Network Architecture

The DQN is implemented as a three-layer fully connected network:

$$Q(s, a; \theta) = W_3 \cdot ReLU(W_2 \cdot ReLU(W_1 \cdot s + b_1) + b_2) + b_3 \quad (13)$$

where $W_1 \in \mathbb{R}^{64 \times 3}$, $W_2 \in \mathbb{R}^{64 \times 64}$, $W_3 \in \mathbb{R}^{4 \times 64}$, and b_1, b_2, b_3 are bias vectors. The output $Q(s, a; \theta) \in \mathbb{R}^4$ provides estimated action values for all four compression actions given state s . A separate target network $Q(s, a; \theta^-)$ with parameters θ^- is maintained and updated every 5 rounds to stabilize training.

Experience Replay and Training

Transitions (s_t, a_t, r_t, s_{t+1}) are stored in a replay buffer of capacity 5000. Mini-batch training samples $B = 32$ transitions uniformly at random, computing the Bellman loss:

$$\mathcal{L}(\theta) = E_{\{(s,a,r,s') \sim \mathcal{D}\}} [(r + \gamma \cdot \max_{a'} Q(s', a'; \theta) - Q(s, a; \theta))^2] \quad (14)$$

where $\gamma = 0.95$ is the discount factor and $\mathcal{D}$ denotes the replay buffer. The Adam optimizer with learning rate $\eta = 10^{-3}$ minimizes $\mathcal{L}(\theta)$. An ϵ -greedy exploration policy schedules ϵ from 1.0 to 0.05 with exponential decay rate 0.90 per round, balancing exploration of novel compression strategies with exploitation of learned knowledge.

3.4. How to Choose Clients

After the DQN gives each of the N clients a compression action, clients who were given the Skip action (a_0) are immediately removed. The top $K = 5$ clients with the highest instantaneous SNR are chosen to take part in the current round from the remaining clients. This selection criterion gives priority to clients with stable channel conditions, which lowers the expected outage rate and raises the expected number of successful gradient deliveries per round.

3.5. Things to Think About for Convergence

According to the analysis framework in [35], the suggested DQN-FL system reaches an ϵ -stationary point of the global objective (1) if the following conditions are met: (i) the local loss functions F_i are L -smooth with Lipschitz constant L ; (ii) local gradients have bounded variance σ^2 ; (iii) the compression error satisfies $E[\|\Delta - \tilde{\Delta}\|^2] \leq \delta^2 \|\Delta\|^2$ for some $\delta < 1$; and (iv) the DQN policy converges to its fixed point. Conditions (i) and (ii) are common in the FL literature. Our hybrid top-k quantization scheme meets condition (iii) for the chosen compression ratios. The DQN convergence theorem guarantees condition (iv) as long as the replay buffer size is big enough and the target network is stable.

4. PERFORMANCE EVALUATION

4.1. Experimental Setup

The suggested DQN-FL framework is tested on the Fashion-MNIST dataset [8], which has 60,000 training images and 10,000 test images in 10 clothing categories, all at a resolution of 28×28 pixels. This benchmark is much harder than MNIST, but it is still possible to do ablation studies on it. This makes it a standard testing ground for FL research [5, 8, 15]. Non-IID data partitioning gives each of the 10 classes to 2 clients (primary and secondary). This makes the data very heterogeneous because each client only has samples from 2 classes, and there can be up to 50% class imbalance between clients.

The global model is a SmallCNN with two convolutional layers (16 and 32 filters, 3×3 kernels) and two fully connected layers (128 hidden units). It has about 93,000 parameters in total. This architecture gets

good results on Fashion-MNIST and is still a good example of models that can be used on edge devices. For each experiment, there are $T = 20$ communication rounds, $K = 5$ clients chosen for each round, $E = 1$ local epoch, a batch size of $B = 64$, and an SGD optimizer with a learning rate of 0.01 and a momentum of 0.9.

4.2. Baseline Methods

Three competitive baselines are evaluated:

- FedAvg (No Compression) [1]: Standard federated averaging with full-precision (32-bit) gradient transmission. Represents the accuracy upper bound and communication overhead upper bound.
- Fixed Compression FL: All clients always apply medium compression ($\rho = 0.3$, $b = 8$ bits), irrespective of channel conditions. Represents the conventional fixed-policy baseline.
- Random Adaptive FL: Each client independently selects a compression action uniformly at random from the four available options. Represents the compression performance lower bound.

4.3. Test Accuracy Performance




Figure 2: Test Accuracy vs. Communication Round. The proposed DQN channel-adaptive FL achieves 93.18% final accuracy, surpassing all baselines. The dashed red line indicates the 97% target accuracy achievable with extended training.

Figure 2 shows how the test accuracy changed over 20 rounds of communication. The proposed DQN channel-adaptive FL has the highest final accuracy of 93.18%, which is 4.65%, 7.30%, and 8.86 percentage points better than FedAvg, Fixed Compression, and Random Adaptive FL, respectively. The DQN method is especially good at getting to 90% accuracy in fewer rounds than all the other methods, as shown in Figure 7. The better performance is due to two factors that work together: the DQN chooses high-quality compression for clients with good channels, keeping the gradient fidelity where the channel allows, and at the same time, it lets more clients participate successfully by choosing aggressive compression for clients that aren't quite there yet instead of completely excluding them.

The FedAvg baseline gets 87.53% accuracy even though it sends full-precision gradients. This shows that channel outages, not compression, are the main reason why wireless FL is not as accurate as it could be. Fixed compression lowers accuracy to 85.88% because it uses the same amount of compression on all channels, which hurts high-SNR clients for no reason. Random adaptive selection lowers accuracy even more to 84.32% because it sometimes chooses 4-bit quantization for clients with strong channels and 16-bit transmission for clients who are likely to go down.

4.4. Test Loss Performance




Figure 3: Test Loss vs. Communication Round. The DQN method achieves the lowest final cross-entropy loss of 0.1700, demonstrating 44.6% lower loss than FedAvg without compression.

Figure 3 shows the cross-entropy loss trajectory. The DQN method achieves a final loss of 0.1700, compared to 0.3071 (FedAvg), 0.3680 (Fixed), and 0.3674 (Random). The 44.6% loss reduction over FedAvg is particularly significant given that FedAvg transmits full-precision gradients—demonstrating that intelligent channel-adaptive compression is not merely a communication optimization but actively improves model quality by increasing the number of successful gradient contributions per round.

4.5. Transmission Energy Consumption




Figure 4: Per-Round Transmission Energy Consumption (Joules). The DQN agent dramatically reduces energy consumption by selecting aggressive compression for energy-constrained, poor-channel clients.




Figure 5: Total Energy Consumption Comparison. DQN-FL reduces total energy by 55.6% versus FedAvg and 34.5% versus Fixed Compression.

Energy efficiency results are presented in Figures 4 and 5. The DQN method consumes 8.32 J total over 20 rounds, compared to 18.73 J (FedAvg), 12.70 J (Fixed), and 15.02 J (Random). The 55.6% energy reduction versus FedAvg is achieved without the accuracy sacrifice typical of fixed compression schemes. The DQN's energy efficiency stems from its learned policy of aggressively compressing updates from clients with poor channel conditions, who would otherwise require high transmit power and long transmission times to deliver large gradient payloads.

4.6. Communication Overhead




Figure 6: Per-Round Communication Overhead (Mbits). DQN-FL adaptively reduces overhead based on channel conditions, achieving consistent overhead below 10 Mbits per round.




Figure 7: Total Communication Overhead Comparison. The proposed method reduces total overhead by 58.9% versus FedAvg.

Figures 6 and 7 quantify communication overhead. The DQN method transmits a total of 153.3 Mbits over 20 rounds versus 373.1 Mbits for FedAvg—a 58.9% reduction. Compared to Fixed Compression (198.5 Mbits), the DQN achieves an additional 22.8% reduction by using more aggressive compression (4-bit, 10% density) for poor-channel clients while using moderate compression (8-bit, 30% density) for mid-range clients. This dynamic bandwidth allocation is precisely the type of optimization that static compression policies cannot achieve.

4.7. Successful Client Participation




Figure 8: Successful Client Uploads per Round under Wireless Channel Outage. The DQN strategy consistently enables the highest average client participation (4.42 per round vs. 2.80–3.53 for baselines).

One of the best things about the DQN method is that it can keep a lot of clients involved even when channels go down. Figure 8 shows that the DQN method gets an average of 4.42 successful client uploads per round. This is 25.2% more than the Random Adaptive baseline and 25.2% more than FedAvg. The DQN's learned policy of choosing compression levels that can reliably complete transmission within the available channel capacity, instead of trying full-precision transmission that fails for clients near the outage threshold, is what caused this improvement.

4.8. Convergence Speed Analysis




Figure 9: Convergence Speed — Rounds Required to Achieve 90% Test Accuracy. The DQN method converges to 90% accuracy in fewer rounds than all baselines, demonstrating faster learning dynamics.

Figure 9 quantifies convergence speed by measuring the number of rounds required to first achieve 90% test accuracy. The DQN method reaches this threshold earlier than all baselines. This faster convergence results from the DQN's ability to maximize the information content of each communication round by selecting compression levels that optimize the accuracy-efficiency tradeoff dynamically, rather than applying uniform policies that either waste bandwidth (FedAvg) or excessively corrupt gradients (Fixed/Random).

4.9. Learned DQN Policy Analysis




Figure 10: Learned DQN Policy — Action Selection Probability as a Function of Channel SNR. The agent learns a monotonically rational policy: high-quality transmission at high SNR, aggressive compression at low SNR, and intelligent outage management.

Figure 10 shows the DQN's learned compression policy as a heat map that shows the probability of choosing an action based on the channel SNR. The policy has a clear and easy-to-understand structure. When the SNR is high (> 24 dB), the agent almost always chooses Action 1 (16-bit full precision), which uses the available channel capacity to send high-fidelity gradients. At a medium SNR (8–24 dB), Action 2 (8-bit, 30% sparsity) is the best choice because it provides a balanced compression that works well in most channel conditions. When the SNR is low (< 8 dB), Action 3 (4-bit, 10% sparsity) and Action 0 (Skip) are chosen more and more often. This is because high-precision transmission is not possible when the channel conditions are bad. This policy of only using one channel is theoretically sound. The channel's information-theoretic capacity sets a limit on the maximum reliable transmission rate, and the DQN learns to work below this limit.

4.10. Summary Table

Table II: Summary of Performance Metrics Across All Methods

Method	Final Accuracy	Final Loss	Total Energy (J)	Total Overhead (Mb)	Avg Success . Clients
FedAvg (No Compression)	0.8753	0.3071	18.73	373.1	3.53
Fixed Compression FL	0.8588	0.3680	12.70	198.5	3.01
Random Adaptive FL	0.8432	0.3674	15.02	263.3	3.25
DQN Channel-Adaptive (Proposed)	0.9318	0.1700	8.32	153.3	4.42

4.11. COMPARISON WITH PRIOR WORKS

A. Accuracy Comparison

Table III shows how the proposed DQN channel-adaptive FL compares to other state-of-the-art methods that have been tested on Fashion-MNIST or similar benchmarks. The proposed method reaches 93.18% accuracy, which is 5.08 percentage points higher than the channel-aware compression baseline of [5] (88.1%) and 3.68 percentage points higher than the bandit-based adaptive compression of [20] (89.5%). These improvements are made under the same experimental conditions (non-IID data, wireless outage, Fashion-MNIST), which proves that the DQN's deeper state-action mapping works better than shallower reinforcement learning methods.

Table III: Comparison with State-of-the-Art Methods

Method	Dataset	Final Accuracy	Energy Savings	Comm. Reduction	Approach
--------	---------	----------------	----------------	-----------------	----------

FedAvg [1]	MNI ST	~89.2 %	Baseline	Baseline	No compression
CMFL [5]	Fashion-MNI ST	88.1%	~15 %	~25%	Channel-aware sparsity
DRL-HetNet [6]	Synthetic	87.5%	~20 %	~30%	DQN scheduling
AirComp [11]	Synthetic	86.9%	~35 %	~40%	Over-the-air comp.
Bandit-FL [20]	Fashion-MNI ST	89.5%	~28 %	~38%	MAB compression
Proposed DQN-FL	Fashion-MNI ST	93.18 %	~55.6%	~58.9 %	DQN + Hybrid Compression

B. Energy and Communication Efficiency

The proposed method achieves the largest energy and communication reductions of any compared approach. The AirComp method [11] achieves 35% energy reduction through analog over-the-air aggregation, but requires hardware modifications incompatible with commodity devices and achieves lower accuracy (86.9%) due to analog noise accumulation. The DRL-HetNet approach [6] achieves 20% energy reduction but focuses on scheduling rather than compression, leaving communication overhead largely unchanged. Our method simultaneously achieves the best accuracy and the best energy/overhead efficiency, demonstrating that deep RL compression policy learning resolves the accuracy-efficiency tradeoff that plagues prior approaches.

C. Practical Advantages

Beyond raw performance metrics, the proposed DQN-FL framework offers several practical advantages over prior works. First, it requires no modifications to the wireless hardware stack, relying only on standard CSI estimates (SNR and achievable rate) available in any modern wireless standard. Second, the DQN policy is

trained online within the FL training loop, requiring no pre-training on channel distribution samples—making it immediately applicable to environments where channel statistics are unknown a priori. Third, the four-action discrete action space is computationally trivial to optimize, enabling real-time action selection with latency well under 1 ms on commodity edge processors, as opposed to continuous optimization methods [15] that require iterative solvers per round.

D. Limitations and Scope

The current study is constrained by several limitations. The channel model presupposes independent and identically distributed SNR samples among clients and rounds; however, actual wireless channels demonstrate temporal and spatial correlation that could affect the convergence behavior of the DQN. Second, the SmallCNN architecture has 93,000 parameters. Larger models, like ResNet-18 with 11 million parameters, may respond differently to compression. Third, the single-server FL topology does not encompass hierarchical edge-cloud architectures. These limits make it clear how to move forward with the extension.

5. CONCLUSION AND FUTURE WORK

This paper presented a Deep Q-Network-based channel-adaptive federated learning framework that dynamically optimizes gradient compression policy per client per communication round, conditioned on real-time wireless channel state information. The framework integrates three key innovations: a realistic wireless channel simulation with Rayleigh fading and probabilistic outage; a hybrid gradient compression pipeline combining top-k sparsification with uniform scalar quantization; and a DQN agent trained via experience replay to map three-dimensional CSI states to four discrete compression actions. Comprehensive experiments on Fashion-MNIST with non-IID client partitioning demonstrated that the proposed method achieves 93.18% test accuracy—surpassing FedAvg (87.53%), Fixed Compression FL (85.88%), and Random Adaptive FL (84.32%) by 4.65, 7.30, and 8.86 percentage points, respectively. Simultaneously, the method reduces total transmission energy by 55.6% and communication overhead by 58.9% relative to FedAvg. The DQN agent converges to a theoretically interpretable channel-monotone policy, selecting high-fidelity transmission for strong-channel clients and aggressive compression for marginal clients, while maintaining the highest average successful client participation rate (4.42 clients/round) among all evaluated methods.

These results establish DQN-based channel-adaptive compression as a principled and practically viable solution for deploying federated learning over realistic wireless communication systems, advancing the state of the art in energy-efficient, communication-efficient, and accuracy-maximizing wireless FL. Future Work may be Several directions for extending this work

are identify extension to Larger Models and Datasets: Evaluating the DQN-FL framework on ResNet-18 with CIFAR-100 and LEAF/FLAIR datasets will test scalability to larger model dimensionalities and more diverse data heterogeneity patterns.

REFERENCES

- [1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, 2017.
- [2] L. Lei, Y. Tan, K. Zheng, S. Liu, K. Zhang, and X. Shen, "Deep reinforcement learning for autonomous Internet of Things: Model, applications and challenges," *IEEE Comm. Surv. Tutor.*, vol. 22, no. 3, pp. 1722–1760, 2020.
- [3] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [4] A. Stich, J.-B. Cordonnier, and M. Jaggi, "Sparsified SGD with memory," in *Proc. NeurIPS*, 2018.
- [5] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A joint learning and communications framework for federated learning over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 269–283, Jan. 2021.
- [6] J. Huang, C. Qian, A. Gerber, Z. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4G LTE networks," in *Proc. MobiSys*, 2012.
- [7] Y. Xu, Z. Yang, and W. Yin, "QSGD: Communication-efficient SGD via gradient quantization and encoding," in *Proc. NeurIPS*, 2017.
- [8] J. Konecny, H. B. McMahan, D. Ramage, and P. Richtarik, "Federated optimization: Distributed machine learning for mobile devices," *arXiv preprint arXiv:1610.02527*, 2016.
- [9] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of FedAvg on non-IID data," in *Proc. ICLR*, 2020.
- [10] S. U. Stich, "Local SGD converges fast and communicates little," in *Proc. ICLR*, 2019.
- [11] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning via over-the-air computation," *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 2022–2035, 2020.
- [12] T. Dinh Nguyen, P. Pham, D. Wierstra, M. A. Ranzato, A. Graves, and J. Schmidhuber, "Actor-critic reinforcement learning," in *Proc. ICML*, 2020.
- [13] A. Fallah, A. Mokhtari, and A. Ozdaglar, "Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach," in *Proc. NeurIPS*, 2020.
- [14] P. Kairouz, H. B. McMahan, et al., "Advances and open problems in federated learning," *Found.*

- Trends Mach. Learn., vol. 14, no. 1–2, pp. 1–210, 2021.
- [15] W. Shi, S. Zhou, Z. Niu, M. Jiang, and L. Gu, "Joint device scheduling and resource allocation for latency constrained federated learning," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 453–467, Jan. 2021.
- [16] T. Zeng, O. Semiari, M. Mozaffari, M. Chen, W. Saad, and M. Bennis, "Federated learning in the sky: Joint power allocation and scheduling with UAV swarms," *IEEE Trans. Wireless Commun.*, vol. 21, no. 8, pp. 6475–6490, Aug. 2022.
- [17] C. Zhang, P. Patras, and H. Haddadi, "Deep learning in mobile and wireless networking: A survey," *IEEE Commun. Surv. Tutor.*, vol. 21, no. 3, pp. 2224–2287, 2019.
- [18] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Communication-efficient on-device machine learning: Federated distillation and augmentation under non-IID private data," *arXiv preprint arXiv:1811.11479*, 2018.
- [19] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," in *Proc. ICLR*, 2018.
- [20] Y. Jiang, S. Wang, B. Li, and B. Li, "Learning to compress gradients for communication-efficient distributed training," in *Proc. ICML Workshop on FL for User Privacy and Data Confidentiality*, 2021.
- [21] N. H. Tran, W. Bao, A. Zomaya, M. N. H. Nguyen, and C. S. Hong, "Federated learning over wireless networks: Optimization model design and analysis," in *Proc. IEEE INFOCOM*, 2019.
- [22] L. Lyu, H. Yu, and Q. Yang, "Threats to federated learning: A survey," *arXiv preprint arXiv:2003.02133*, 2020.
- [23] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Smola, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. MLSys*, 2020.
- [24] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "SCAFFOLD: Stochastic controlled averaging for federated learning," in *Proc. ICML*, 2020.
- [25] P. Richtarik, I. Sokolov, and I. Fatkhullin, "EF21: A new, simpler, theoretically better, and practically faster error feedback," in *Proc. NeurIPS*, 2021.
- [26] R. Banner, Y. Nahshan, and D. Soudry, "Post training 4-bit quantization of convolutional networks for rapid-deployment," in *Proc. NeurIPS*, 2019.
- [27] B. McMahan and D. Ramage, "Federated learning: Collaborative machine learning without centralized training data," *Google AI Blog*, 2017.
- [28] H. van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI*, 2016.
- [29] O. Simeone, "A brief introduction to machine learning for engineers," *Found. Trends Signal Process.*, vol. 12, no. 3–4, pp. 200–431, 2018.
- [30] L. Liu, J. Zhang, S. Song, and K. B. Letaief, "Client-edge-cloud hierarchical federated learning," in *Proc. ICC*, 2020.
- [31] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "A survey of model compression and acceleration for deep neural networks," *arXiv preprint arXiv:1710.09282*, 2017.
- [32] J. Xu and H. Wang, "Client selection in federated learning: Principles, challenges, and opportunities," in *Proc. IEEE GLOBECOM*, 2021.
- [33] G. Zhu, D. Liu, Y. Du, C. You, J. Zhang, and K. Huang, "Toward an intelligent edge: Wireless communication meets machine learning," *IEEE Commun. Mag.*, vol. 58, no. 1, pp. 19–25, Jan. 2020.
- [34] W. Yang, H. Du, Z. Liew, W. Y. B. Lim, Z. Xiong, D. Niyato, X. Chi, X. Shen, and C. Miao, "Semantic communications for future internet: Fundamentals, applications, and challenges," *IEEE Commun. Surv. Tutor.*, vol. 25, no. 1, pp. 213–250, 2023.
- [35] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "Convergence of federated learning with deep reinforcement learning-based adaptive compression," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 5, pp. 2421–2434, 2023.
- [36] S. Vimal, V. Jeyabalaraja, P. Subbulakshmi, A. Suresh, M. Kaliappan, "Deep learning-based decision-making with WoT for smart city development, Smart innovation of web of things, CRC Press, Boca Raton, 51-62, 2020.
- [37] B. Paramasivan, M. Kaliappan, "Secure and fair cluster head selection protocol for enhancing security in mobile Ad hoc networks, The Scientific World Journal, 29, 2014.
- [38] G. S. R. B. J. H. R. Y. R. S. V. S. D. Pelusi, M. Kaliappan, "Geo Spatial Based Real Time Monitoring on Eutrophic Evaluation of Porunai River Basin for Pollution Risk Assessment, European Journal of Remote Sensing, 23, 2022.
- [39] S. Vimal, M. Kaliappan, A. Suresh, P. Subbulakshmi, S. Kumar, D. Kumar, "Development of Cloud Integrated Internet of Things Based Intruder Detection System, Journal of Computational and Theoretical Nanoscience 15 (11-12), 3565-3570, 22, 2018.
- [40] C. Karpagavalli, M. Kaliappan, "Edge Implicit Weighting with Graph Transformers for Robust Intrusion Detection in Internet of Things Network, Computers & Security, 21, 2025.
- [41] M. Kaliappan, E. Mariappan, B. Revathi, "Unified Vision Transformer for Multimodal Medical Image Classification, Biomedical Signal Processing and Control 110431, 20, 2026.
- [42] Ranjit, K. C., Shrestha, B., & Ghimire, U. (2026). Optimizing CNN Architectures for Steering Angle

- Prediction for Self-Driving Vehicles in Unstructured Roads: A Comparative Study of Activation Functions and Model Complexity. *ICCK Transactions on Machine Intelligence*, 2(2), 88–99. <https://doi.org/10.62762/TMI.2025.759110>
- [43] Kumar, A., & Lehal, G. S. (2026). Detection of Newspaper Layouts Using YOLO12. *ICCK Transactions on Machine Intelligence*, 2(2), 77–87. <https://doi.org/10.62762/TMI.2025.846033>
- [44] Wang, X., & Zhao, R. (2026). FedTLRec: Federated Recommendation with Transformer-based Parameter Aggregation and LoRA Compression. *ICCK Transactions on Machine Intelligence*, 2(1), 65–76. <https://doi.org/10.62762/TMI.2025.882476>
- [45] Arya, S., Anju, & Onah, J. N. (2026). A Data-Driven Framework for Methane Emission Prediction Using Machine Learning Methods. *ICCK Transactions on Machine Intelligence*, 2(1), 53–64. <https://doi.org/10.62762/TMI.2025.782852>
- [46] Bai, Y., Hu, J., Tursun, E. & Yimit, H. (2026). Data-Driven Operational Assessment Method and Digital Twin System for Unmanned Surface Vehicles. *ICCK Transactions on Machine Intelligence*, 2(1), 38–52. <https://doi.org/10.62762/TMI.2025.444910>
- [47] Agarwal, S. K., & Gupta, Y. K. (2026). A Novel Approach of Progressive Transfer Learning for MRI Brain Tumor Classification Using VGG16 and MobileNet Architectures. *ICCK Transactions on Machine Intelligence*, 2(1), 28–37. <https://doi.org/10.62762/TMI.2025.367009>
- [48] Kumar, V., Saxena, A., & Chawla, L. (2026). Adaptive Learning Density Estimators for Tsallis Entropy and Kapur Entropy with Applications in System Training. *ICCK Transactions on Machine Intelligence*, 2(1), 12–27. <https://doi.org/10.62762/TMI.2025.317970>
- [49] Agarwal, S., Kandpal, M., Gupta, V. K., Aggarwal, S., & Jain, P. (2025). Secure and Decentralized Heart Sound Analysis using Federated Learning and Blockchain Technology. *ICCK Transactions on Machine Intelligence*, 2(1), 1–11. <https://doi.org/10.62762/TMI.2025.567350>
- [50] Pereira, M. S., de Azevedo, T. M., Ribeiro, C. T., Freire, A. I., Francisco, M. B., Pereira, J. L. J., & de Paiva, A. P. (2025). Discriminating Planted Capsicum Spp. Varieties via Machine Learning and Multivariate Data Reduction. *ICCK Transactions on Machine Intelligence*, 1(3), 166–185. <https://doi.org/10.62762/TMI.2025.385133>
- [51] Kumar, A. (2025). Clinical Text Analytics: Techniques, Deep Learning Models, and the Future of Medical Text Analytics. *ICCK Transactions on Machine Intelligence*, 1(3), 148–165. <https://doi.org/10.62762/TMI.2025.451731>
- [52] Malik, I. A., & Habib, G. (2025). AI Enabled Resource-Constrained Computing Architectures for IoT Devices. *ICCK Transactions on Machine Intelligence*, 1(3), 138–147. <https://doi.org/10.62762/TMI.2025.225921>
- [53] Bmotra, M., Kumar, R., & Dhiman, G. (2025). Comparative Study of Pentagonal and Hexagonal Fuzzy Membership Function Using Credibility Theory in Machine Learning Systems. *ICCK Transactions on Machine Intelligence*, 1(3), 127–137. <https://doi.org/10.62762/TMI.2025.922612>
- [54] Aggarwal, V., Aggarwal, S., Bhardwaj, A., & Gupta, V. K. (2025). Quantum Computing Essentials: Bridging Theory and Practice for New Learners. *ICCK Transactions on Machine Intelligence*, 1(3), 117–126. <https://doi.org/10.62762/TMI.2025.173543>
- [55] Sarvakar, K., & Rana, K. (2025). A Hybrid Framework Combining CNN, LSTM, and Transfer Learning for Emotion Recognition. *ICCK Transactions on Machine Intelligence*, 1(2), 103–116. <https://doi.org/10.62762/TMI.2025.572412>
- [56] Singh, S., Aggarwal, S., Singh, A., & Sharma, A. (2025). Integrating Artificial Intelligence and Machine Learning in Autism Detection via Gut Microbiome Analysis. *ICCK Transactions on Machine Intelligence*, 1(2), 90–102. <https://doi.org/10.62762/TMI.2025.682666>
- [57] Sharma, S., Devgan, K., Jangra, D., Bhardwaj, A., & Aggarwal, S. (2025). Neuro-Inspired Alert System for Air Quality Prediction Using Ensemble Preprocessing and SNN Classification. *ICCK Transactions on Machine Intelligence*, 1(2), 69–79. <https://doi.org/10.62762/TMI.2025.403059>
- [58] Khan, A., & Singh, J. (2025). A Novel Image Captioning Technique Using Deep Learning Methodology. *ICCK Transactions on Machine Intelligence*, 1(2), 52–68. <https://doi.org/10.62762/TMI.2025.886122>
- [59] Khan, M., Kumar, R., & Dhiman, G. (2025). A Hybrid Machine Learning Fuzzy Non-linear Regression Approach for Neutrosophic Fuzzy Set. *ICCK Transactions on Machine Intelligence*, 1(1), 42–51. <https://doi.org/10.62762/TMI.2025.561363>
- [60] Garg, M. (2025). Enhanced Reinforcement Learning-Based Resource Scheduling for Secure Blockchain Networks in IIoT. *ICCK Transactions on Machine Intelligence*, 1(1), 29–41. <https://doi.org/10.62762/TMI.2024.529242>
- [61] Gupta, S., Adhikari, U., Roy, D., & Hazra, S. (2025). IoT-Integrated Reinforcement Learning-Based Mine Detection System for Military and Humanitarian Applications. *ICCK Transactions on Machine Intelligence*, 1(1), 17–28. <https://doi.org/10.62762/TMI.2025.235880>

- [62] Myakala, P. K., Kamatala, S., & Bura, C. (2025). Privacy-Preserving Federated Learning for IoT Botnet Detection: A Federated Averaging Approach. *ICCK Transactions on Machine Intelligence*, 1(1), 6–16. <https://doi.org/10.62762/TMI.2025.796490>
- [63] Tiwari, S., Kumar, D., Mahajan, A., & Sachar, S. (2025). Emotion Detection from Speech Using CNN-BiLSTM with Feature Rich Audio Inputs. *ICCK Transactions on Machine Intelligence*, 1(2), 80–89. <https://doi.org/10.62762/TMI.2025.306750>