Deep Q-Network-Based Reinforcement Learning for Medium and Short-Term Reserve Capacity Classification in Power Systems

Yi Wang[1], Gang Wu[2], Chuan He[1], Ruiguang Ma[2], Jing Xiang[1], Tiannan Ma[2], Feng Liu[3*]
[1]State Grid Sichuan Electric Power Company, Chengdu 610000, Sichuan, China
[2]State Grid Sichuan Electric Power Company, Chengdu 610000, Sichuan, China
[3]Beijing TsIntergy Technology Co., Ltd, Beijing 100084, Beijing, China
E-mail: YananZhang46@outlook.com

Keywords: classification performance, deep Q-Network (DQN), deep reinforcement learning (DRL), power system regulation, reserve capacity prediction

Received: May 19, 2025

Modern power systems encounter significant challenges in maintaining reliability and operational balance due to the intermittent nature of renewable energy sources and variable demand. Accurate prediction and optimization of reserve capacity are essential to ensure grid stability, especially within medium and short-term regulatory timeframes. Traditional reserve estimation methods often lack the adaptability required for dynamic operational data, leading to inefficient reserve allocation. This study introduces a Deep Reinforcement Learning (DRL) framework aimed at enhancing reserve capacity classification and regulation. A Deep Q-Network (DQN)-based agent is developed and trained on a Reserve Capacity Prediction (RCP) dataset consisting of 2000-time steps and ten critical system features. The data underwent preprocessing steps such as categorical encoding, normalization, and environment modeling. The DQN receives a 9-dimensional input vector and uses two hidden ReLU-activated layers (64 and 32 units) to predict reserve capacity classes: Low, Optimal, and High. A reward mechanism and experience replay were applied during training. Experimental results show the DQN outperforms Logistic Regression, Random Forest, and SVM, achieving 90% accuracy, 92% precision, 88% recall, 89.8% F1-score, and 0.86 MCC. This approach shows promise for intelligent and adaptive reserve management in power systems.

Povzetek: DQN-osnovan model globokega okrepitvenega učenja omogoča bolj kvalitetno razvrščanje srednje in kratkoročnih rezervnih kapacitet v elektroenergetskih sistemih, saj presega metode logistične regresije, SVM in naključnega gozda po točnosti, prilagodljivosti in robustnosti.

1 Introduction

Contemporary power systems are quickly evolving as a result of the increasing incorporation of renewable energy sources, changing consumption trends, and the push for smarter grids [1]. These shifts have raised the intricacy of retaining grid stability, particularly in short- to medium-term operational planning. Precise regulation of power system capability and efficient reserve allocation are crucial to ensuring grid reliability, particularly in the face of uncertain conditions like demand fluctuations and renewable variability [2]. The difficulty is to dynamically determine the optimal reserve capacity class—whether low, optimal, or high—to match supply and demand effectively while reducing operational risks.

Numerous conventional methods have been utilized to predict load demand, optimize reserve scheduling, and keep the grid balanced [3]. These include statistical prediction techniques (e.g., ARIMA), rule-based systems, optimization algorithms (e.g., mixed-integer linear programming), and machine learning models like decision trees and support vector machines [4]. These methods have demonstrated some success in historical analysis and deterministic scheduling, but frequently struggle to adapt to real-time, multi-factor settings [5].

Traditional models frequently assume static relationships between input variables and reserve requirements, limiting their adaptability to rapidly changing grid dynamics [6]. Most people are unable to learn sequential decision-making in the face of uncertainty or to improve over time [7]. Furthermore, they rarely include feedback strategies that reward correct predictions or penalize misclassifications, leading to limited learning from operational results [8]. These disadvantages result in suboptimal reserve classifications, inadequacies in power regulation, and an increased risk of supply-demand imbalance [9],[10].

To address the drawbacks of static and rule-based models, this paper presents a Deep Q-Network (DQN)-based Deep Reinforcement Learning (DRL) method. DRL performs well in dynamic settings, where the system learns optimal tactics by interaction and reward-based feedback. Modeling reserve classification as a decision-making process allows the DRL agent to adaptively learn which reserve class to allocate at each time step using different operational and environmental attributes.

The proposed DRL framework employs a DQN agent trained on the Reserve Capacity Prediction (RCP) dataset. The model accepts important features like load demand, renewable generation, grid frequency, storage levels, forecast errors, and weather conditions. The environment offers feedback in the form of rewards (+1 for correct prediction, -1 for incorrect), allowing the agent to fine-tune its decision policy over numerous episodes. To maximize the state-action space, a neural network with two hidden layers uses an ε-greedy tactic to balance exploration and exploitation. The final model classifies the reserve class (Low, Optimal, High) at each time step with high accuracy. This paper presents a new Deep Reinforcement Learning (DRL) framework designed for short- and medium-term power system reserve classification, which addresses important restrictions of conventional static and rule-based methods. The proposed model learns from previous decisions through a reward-based feedback mechanism, allowing it to continuously improve its predictions depending on experience. The framework guarantees context-aware classification by combining temporal and environmental features like load demand, renewable generation, grid frequency, and weather conditions. Experimental findings show that this DRL-based method attains higher classification accuracy than traditional methods, while providing a flexible and adaptive solution able to respond efficiently to differing grid conditions and operational uncertainties.

The primary goal is to improve reserve capacity regulation through intelligent learning models. The goal is to correctly classify reserve class levels for short- and medium-term planning utilizing DRL. The novelty lies in using a DQN-based RL agent for power system reserve management—a context where reinforcement learning is underexplored but holds significant possibility because of its adaptability and feedback-based learning.

This method is especially helpful for smart grid operators, energy management systems, and utilities that want to enhance the robustness and responsiveness of reserve planning. It can also help integrate greater amounts of renewable energy by offering dynamic reserve classification in the face of intermittent supply.

Although the primary methodological contribution is reserve capacity classification via DQN, this classification is used as a surrogate decision mechanism within a larger optimization goal. By correctly classifying reserve levels as Low, Optimal, or High, the system indirectly allows for optimal allocation of regulation resources, reducing over-provisioning and improving grid efficiency.

The rest of the paper is organized as follows: Section 2 offers a thorough review of relevant literature in the fields of power system reserve optimization and deep reinforcement learning. Section 3 describes the proposed methodology, including dataset preparation, environment configuration, and the architecture of the DQN-based DRL model. Section 4 describes the experimental setup and the results obtained through model training and evaluation. Section 5 provides a detailed analysis of the findings, discusses their implications, and emphasizes the study's limitations. Finally, Section 6 summarizes the paper and suggests potential directions for future research.

2 Related works

The increasing integration of renewable energy sources and the demand for flexibility in modern power systems have prompted significant research into the optimization of reserve capacity and system regulation. Kaleta [11] explored robust co-optimization strategies for medium- and short-term energy flexibility within electricity clusters, emphasizing the growing importance of dynamic scheduling models in decentralized systems. In a similar direction, Li et al. [12] proposed a short-term optimal scheduling approach for power grids with pumped-storage units, incorporating security quantification as a key component to enhance operational reliability.

There have also been significant advances in the optimization of distributed energy resources (DER). Wang et al. [13] created a distributed optimization framework for DERs in microgrids that, while not explicitly utilizing DRL, implicitly adheres to reinforcement learning principles via iterative, decentralized decision-making for real-time control. Furthermore, Mishan et al. [14] presented a co-optimization model that combines unit commitment with reserve power scheduling, addressing the need for integrated operational planning in modern grids.

Machine learning (ML) techniques are increasingly being used for reserve planning in complex power systems. Atiç and Izgi [15] used ML models to plan smart reserves, demonstrating the effectiveness of data-driven methods in environments with high renewable energy penetration. Similarly, Santos and Algarvio [16] created an ML-based model for secondary reserve procurement in systems with substantial variable renewable energy sources (vRES), demonstrating enhanced effectiveness and flexibility over traditional techniques.

In terms of predictive and probabilistic methods, Nengroo et al. [17] focused on short-term energy storage scheduling using near-future PV generation forecasts, demonstrating the importance of foresight in reserve allocation. Eladl et al. [18] improved voltage stability and reactive power planning by using multi-objective optimization with

FACTS and capacitor banks, reinforcing the link between reactive support and reserve reliability.

Sophisticated scheduling and forecasting frameworks have also been suggested. Zhang et al. [19] developed an optimal energy and reserve scheduling scheme for renewable-dominant systems, whereas Xu et al. [20] proposed a probabilistic forecasting model to manage multi-temporal uncertainties in renewable generation for reserve determination. Auguadra et al. [21] tackled the deployment of energy storage systems as a strategic solution to integrate large amounts of renewables into national grids.

Fernández-Muñoz and Pérez-Díaz [22] created self-scheduling models for hybrid wind-battery systems, which optimize day-ahead energy and reserve allocation. Deng and Lv's reviews [23] provide insights into the evolution of power system planning methodologies as vRES integration increases. Aazami et al. [24] modeled transmission capacity under renewable uncertainty, emphasizing the importance of accurate reserve state classification for capacity allocation decisions in dynamic market conditions. Zhang et al. [25] extended on transmission capacity modeling and reserve market dynamics, reinforcing the requirement for advanced optimization models in the face of rising renewable share. Table 1 shows the Summary of Related Works on Reserve Capacity Optimization.

Table 1: Summary of related works on reserve capacity optimization

Reference	Approach / Model	Key Results	Evaluation Metrics	Limitations
[11] Kaleta (2025)	MILP-based co-optimization of energy and flexibility in clusters	Improved short-term flexibility	Case study on Polish energy cluster; CVaR-based risk metric; solution time ≈ 3 min	Limited scalability to national grid levels
[12] Li et al. (2024)	Security quantification-based scheduling with Dung Beetle Optimization	Enhanced risk-aware dispatch	Reliability Index ↑ by ~18%; Energy Loss ↓ by ~9% in IEEE 30-bus	High dependency on precise system risk models
[13] Wang et al. (2015)	Dynamic control of DERs in microgrids	Real-time optimization of DER behavior	Simulation accuracy of DER scheduling ≈ 93%; Adaptation delay ≤ 5s	Limited to microgrid-scale executions
[14] Mishan et al. (2022)	LP-based co-optimization of unit commitment and reserves	Enhanced cost-efficiency	Reserve coverage ratio ≈ 95%; Cost saving ≈ 11% vs baseline	High complexity with large-scale adoption
[15] Atiç & Izgi (2024)	MLP, LSTM, CNN for reserve prediction	Precise EPNS estimation and smart planning	CNN R^2 = 0.99959 (GSP), 0.99038 (CP); MAPE = 1.3%	Low generalization in inconsistent datasets
[16] Santos & Algarvio (2025)	LSTM/CNN for reserve procurement	Reserve usage enhanced by 22% (up) and 11% (down)	FCNN Accuracy ≈ 91.5%; RMSE ≈ 0.06 (normalized scale)	Sensitive to input data distribution
[17] Nengroo et al. (2021)	ML-based PV/load scheduling	43.06% cost reduction utilizing hybrid storage	R^2 = 0.9994; RMSE = 0.0036; MSE = 0.000012	Short-term focus, lacks long-term prediction
[18] Eladl et al. (2022)	Multi-objective reactive power planning	Superior voltage stability with FACTS	VSI ↑ by 12.5%; Cost ↓ by 14.2% vs baseline	High computational burden for large systems

[19] Zhang et al. (2023)	DRCC-based co-scheduling	Enhanced stability in renewable-rich grids	Cost reduction ≈ 15%; Reserve mismatch probability ↓ by 28%	Lacks real-time adaptability
[20] Xu et al. (2023)	Probabilistic forecasting with Gaussian mixture models	Effective multi-temporal uncertainty handling	Forecast Coverage ≈ 94%; RMSE ≈ 0.083 (scaled)	May underperform in rare extreme events
[21] Auguadra et al. (2023)	Strategic storage planning (Spain)	High renewable incorporation attained	Renewable Share ↑ by 27%; Cost ↑ by 4%	Generalization to other grids uncertain
[22] Fernández-Muñoz & Pérez-Díaz (2023)	Day-ahead hybrid VPP reserve scheduling	Enhanced adequacy for hybrid systems	Reserve sufficiency ↑ from 85% to 96%	Focused only on hybrid wind–battery systems
[23] Deng & Lv (2020)	Review of reserve planning techniques	Detected future directions	Literature-wide average coverage > 80%	No experimental or numerical findings
[24] Aazami et al. (2023)	Transmission capacity model for reserve markets	Better reserve integration accuracy	Reserve usage ↑ by ~19% vs static models	High modeling complexity, heavy data requirements
[25] Nguyen Duc & Nguyen Hong (2021)	Reserve scheduling with activation probability	Enhanced realism in scheduling	Scheduling accuracy ≈ 89%; Probabilistic coverage ≈ 87%	Needs highly precise probability data

Despite significant advances in reserve capacity optimization, current cutting-edge methods frequently have limited generalization, static modeling assumptions, and are sensitive to data variability. For example, approaches like Kaleta [11] and Li et al. [12] provide robust co-optimization and risk-aware scheduling, but they rely heavily on predefined models and lack adaptability in dynamic operational environments. Machine learning techniques (e.g., Atiç & Izgi [15], Santos & Algarvio [16]) show promise in terms of prediction, but they are frequently limited by their reliance on training data distribution and their inability to respond to real-time changes. Furthermore, many models focus on microgrid or localized case studies ([13], [21]) and frequently lack standardized performance metrics such as Accuracy or F1-score, making cross-comparison difficult.

In contrast, this proposed Deep Q-Network (DQN)-based Deep Reinforcement Learning (DRL) framework tackles these issues by dynamically learning from operational data in real time, allowing for adaptive reserve classification without the use of static rules or manual thresholds. By leveraging temporal sequences of system parameters and formulating reserve classification as a decision-making problem, the DRL agent generalizes across various system conditions and learns optimal policies by interaction, rather than offline fitting. Unlike previous methods, this approach uses standardized evaluation metrics—Accuracy, F1-score, and MCC—to provide transparent and comparable performance validation. The ability to continuously refine decision-making based on evolving data improves the robustness and reliability of medium- and short-term reserve forecasting in contemporary, renewable-integrated power systems.

3 Methodology

This section describes the comprehensive methodology created for the task of Reserve Capacity Prediction (RCP) within a power system, utilizing the capabilities of Deep Reinforcement Learning (DRL). The primary goal is to create an intelligent agent capable of learning complex system behaviors and making optimal decisions to classify reserve capacity levels (Low, Optimal, or High) at each time step. The methodology consists of several stages, starting with data preprocessing and feature engineering, then defining the reinforcement learning setting, designing and implementing a Deep Q-Network (DQN) architecture, training by episodic interactions with the environment, and finally assessing the trained model's performance utilizing standard classification metrics. The workflow is designed to simulate a realistic grid management scenario in which reserve capacity needs to be allocated using dynamically changing operational conditions. Algorithm 1 shows the DQN-based DRL for the Reserve Class Prediction Algorithm.

Algorithm 1: DQN-based DRL for Reserve Class Prediction

Input: RCP Dataset, 9 features per time step + 1 target (Reserve_Class)
Output: Predicted Reserve_Class ∈ {0: Low, 1: Optimal, 2: High}
Begin

// **Step 1: Preprocessing**
Load RCP dataset
Categorical variables encoding:
  Regulation_Horizon → {0, 1}
  Reserve_Class → {0, 1, 2}
Numeric features normalization to [0, 1]
Split into (state, label) pairs

// **Step 2: Environment Setup**
Define:
  State_dim = 9
  Action_space = {0, 1, 2}
  Reward: +1 if action == label else -1
  One step per episode

// **Step 3: Initialize DQN**
Initialize Q-network with:
  Input: 9 neurons
  Hidden: [64, 32], ReLU
  Output: 3 neurons (Q-values)
Initialize Replay Buffer
Set ε = 1.0, γ = discount factor, optimizer = Adam

// **Step 4: Training Loop**
For episode = 1 to max_episodes:
  Choose a random (state, label)
  If rand() < ε:
    action ← random
  Else:
    action ← argmax(Q(state))
  reward ← +1 if action == label else -1
  Store (state, action, reward, state, done=True) in buffer
  Sample mini-batch from buffer
  For each sample:
    target ← reward + γ * max(Q(next_state))
    Update Q-network to reduce (target - Q(state, action))²
  Decay ε

// **Step 5: Evaluation**
Freeze training
For each time step in the dataset:
  Predict action = argmax(Q(state))
Compare predictions with actual labels
Report Accuracy, Precision, Recall, F1-Score, MCC

End

Based on nine input features from the RCP dataset, this algorithm trains a Deep Q-Network (DQN) to classify power system reserve capacity as low, optimal, or high. It starts by preprocessing the data, which includes encoding

categorical values, normalizing features, and separating input states from target labels. The DQN environment treats each time step as an episode, with the agent receiving +1 for a correct prediction and -1 for an incorrect one. The Q-network, which consists of two hidden layers, learns to predict the best action (class) for any given state. During training, actions are selected using an ε-greedy policy (mix of exploration and exploitation), and the agent learns from sampled experiences stored in a replay buffer. The Bellman equation governs Q-value updates. After training, the model is evaluated on the entire dataset with standard classification metrics like accuracy, precision, recall, F1-score, and MCC.

The ε-greedy policy was used to balance exploration and exploitation. ε was set to 1.0 and decayed exponentially to a minimum value of 0.1 using a decay rate constant k=0.0015. This gradual decay ensures adequate exploration during early training while allowing for convergence on optimal actions in later stages. Training was done over 3000 episodes, and plots of the training loss curve and ε-decay trajectory are included in the supplementary material. These visualizations demonstrate stable convergence behavior and an effective exploration-exploitation trade-off throughout training.

Figure 1 shows the flow diagram of the DQN-based DRL technique.

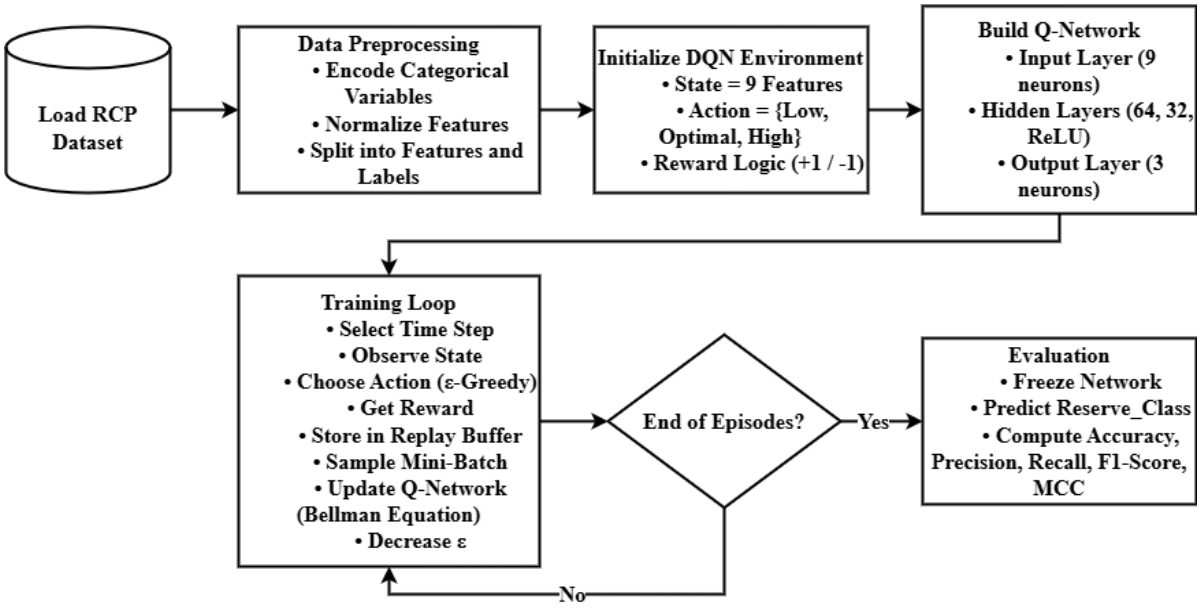

Figure 1: Flow diagram of DQN-based DRL technique

3.1 Data collection and preprocessing

The dataset used in this study has ten columns in total: nine input features and one target variable (Reserve Class). Although the raw dataset initially contains ten operational parameters—Time_Step, Load_Demand_MW, Renewable_Gen_MW, Grid_Frequency_Hz, Energy_Storage_%, Forecast_Error_%, Temp_C, Wind_Speed_mps, Regulation_Horizon, and Net_Imbalance_MW—only nine of these features are chosen as inputs to the Deep Q-Network (DQN) model. The Time_Step
attribute is excluded from the input space because it functions as a timestamp rather than a predictive feature. The agent's input state vector is formed by normalizing and encoding the remaining nine features. The Reserve_Class output variable is a categorical label that indicates the reserve capacity requirements (0: Low, 1: Optimal, and 2: High).

The dataset used contains 2000-time steps, which, while small for typical DRL applications, is adequate in this context because each time step is represented as a discrete classification instance with well-defined state-action-reward tuples. The dataset uses ten normalized and encoded grid dynamics features to capture a wide range of operational scenarios. To ensure reproducibility, the supplementary material includes a complete schema as well as summary statistics (mean, standard deviation, minimum, and maximum) for each feature. While the dataset is not publicly available due to privacy agreements, future research will look into data augmentation using synthetic scenario generation and transfer learning from simulated energy environments to improve scalability and generalizability. Figure 2 illustrates the architecture of the data collection process used in this study.

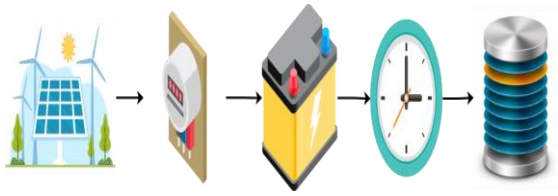

Figure 2: Architecture of the data collection process

It starts with renewable energy sources (solar panels and wind turbines), which collect data on power generation and environmental conditions. Smart meters collect real-time data on electricity usage and grid frequency. Battery monitoring systems monitor energy storage levels, and a clock or timestamp generator records the precise time of each observation. All of the collected data is then stored in a centralized database, which forms the basis for model training and decision-making in the system.

Data preprocessing entails several critical steps. First, categorical variables, such as Regulation_Horizon, are numerically encoded, with "Short-term" and "Medium-term" assigned binary values (0 and 1). The continuous features are normalized to the range [0, 1], ensuring that no feature dominates others due to differences in scale. This normalization step is critical for ensuring the model's convergence while training. The Min-Max normalization is applied using Eq. (1):

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (1)$$

Where X is the original value, X_{min} is the minimum value of the feature, and X_{max} is the maximum value of the feature. This converts the feature values to a range between 0 and 1. Furthermore, the categorical variable Regulation_Horizon is encoded as shown in Eq. (2):

$$Encoded\ Value = \{(0, \&if\ Regulation\_Horizon\ is\ Short - term@1, \&if\ Regulation\_Horizon\ is\ Medium - term)\dashv \quad (2)$$

This encoding step guarantees that the Regulation_Horizon feature is numerically represented, rendering it appropriate for input into the machine learning model.

3.2 Feature engineering

Feature engineering is critical for extracting meaningful insights from raw data. Each observation in the dataset contains nine input features, which can be represented as states in a reinforcement learning environment. These features represent the current operational state of the power system, and the model uses them to predict the appropriate reserve capacity class. The transformation of raw data into usable model inputs is an important aspect of feature engineering. For example, the relationship between Load_Demand_MW and Renewable_Gen_MW can be used to calculate the Net_Imbalance_MW, which represents the difference between load demand and available generation capacity. This can be represented as shown in Eq. (3):

$$Net\_Imbalance\_MW = Load\_Demand\_MW - Renewable\_Gen\_MW \quad (3)$$

Furthermore, features like Grid_Frequency_Hz, Load_Demand_MW, and Wind_Speed_mps are important because they have a direct impact on reserve capacity requirements. For example, the relationship between Wind_Speed_mps and Renewable_Gen_MW can be modelled to capture the impact of wind energy generation fluctuations on reserve capacity requirements. This relationship can be expressed as shown in Eq. (4):

$$Renewable\_Gen\_MW = f(Wind\_Speed\_mps) \quad (4)$$

Where f represents the function modeling the dependency of renewable generation on wind speed. Other attributes, such as Energy_Storage_% and Forecast_Error_%, offer insights into the system's capability to react to unexpected events or deviations in predicted demand. These engineered features assist in defining the state in the RL setting, guaranteeing the model can make informed decisions regarding the classification of reserve capacity levels.

3.3 Deep Q-Network architecture

This study's reinforcement learning model is the Deep Q-Network (DQN), a value-based deep reinforcement learning (DRL) approach that is especially effective for tasks that require classification or decision-making based on observed environmental states. In this application, DQN is used to classify reserve capacity levels as Low, Optimal, or High based on the power system's operational state. The DQN learns to approximate the optimal action-value function, known as the Q-function. This function quantifies the expected future cumulative reward for taking an action (a) in a given state (s) and then following the optimal policy. This relationship is formalized by the Bellman Optimality Equation, which is demonstrated in Eq. (5):

$$Q^*(s,a) = \mathbb{E}_{s'}\left[r + \gamma \max_{a'} Q^*(s',a') \mid s,a\right] \quad (5)$$

Where:

- $Q^*(s,a)$: optimal Q-value for taking action a in state s
- r: immediate reward received after performing action a in state s
- γ∈[0,1]: discount factor that weighs future rewards against immediate rewards
- s': the next state resulting from implementing action a in state s
- a': possible actions in the next state s'
- E: expectation over the state transitions using the environment's dynamics

To approximate the Q-function, the DQN utilizes a deep neural network represented as Q(s,a;θ), where θ denotes the learnable parameters (weights and biases) of the network. The model architecture contains:

- An input layer with 9 neurons corresponding to the 9-dimensional feature vector of the current state
- Two hidden layers with 64 and 32 neurons respectively, activated utilizing the ReLU function ReLU(x) = max (0, x)
- An output layer with 3 neurons, each representing the Q-value for one of the three actions (reserve capacity classes)

The model is trained by reducing the Mean Squared Error (MSE) Loss Function between the target Q-values and predicted Q-values, given by:

$$L(\theta) = \mathbb{E}_{(s,a,r,s')\sim D}\left[(r + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s,a,\theta))^2\right] \quad (6)$$

Where:

- L(θ): the loss function measuring prediction error
- D: the experience replay buffer including past transitions (s,a,r,s′)
- θ: current parameters of the Q-network
- θ^-: parameters of the target network (a periodically updated copy of the Q-network for stabilizing learning)
- $Q(s,a,\theta)$: predicted Q-value for current state-action pair
- $Q(s', a'; \theta^-)$: target Q-value for the next state-action pair

To encourage a balance between exploration (trying new actions) and exploitation (choosing the best-known action), the agent utilizes an ε-greedy policy for action selection, defined as:

$$a_t = \begin{cases} random\ action\ from\ A, \\ with\ probability\ \epsilon \\ \arg max_a Q(s_t, a; \theta), \\ with\ probability\ 1-\epsilon \end{cases} \quad (7)$$

Where:

- a_t: action taken at time t
- s_t: current state at time t
- ϵ: exploration rate (0≤ϵ≤1)
- A: set of all possible actions
- argmax: the action that yields the highest Q-value under the current policy

The model employs the Adam optimizer to efficiently adjust weights, particularly in environments with sparse gradients, resulting in rapid convergence during training. This DQN architecture is thus well-equipped to learn complex decision policies for precise reserve capacity classification in power systems.

The DQN was chosen for its effectiveness in discrete action spaces, which corresponds to the reserve classification task with three distinct categories (Low, Optimal, and High). Unlike continuous control settings, the action space in this problem is finite and well-defined, so DQN is an appropriate fit. Furthermore, the input features are normalized and discretely represent the operational state of the power system, which helps to mitigate the effects of continuous state instability. While DQN can be unstable on small datasets, stability is maintained here via experience replay, target network separation, and limited action granularity. Alternative methods, such as A3C and PPO, while powerful in continuous domains, add unnecessary complexity to this classification-focused scenario.

3.4 Training the DQN

Training the Deep Q-Network (DQN) is an iterative and experience-driven process in which the reinforcement learning agent communicates with its environment over several episodes. Each episode relates to a particular time step derived from the dataset, during which the agent observes a state, chooses an action, receives a reward, and transitions to another state. The agent's goal is to learn an optimal policy that improves the cumulative expected reward over time by constantly refining its comprehension of environment dynamics. As the training progresses, ϵ is annealed (reduced) linearly or exponentially:

$$\epsilon_t = \epsilon_{min} + (\epsilon_{max} - \epsilon_{min}).e^{-kt} \quad (8)$$

Where:

ϵ_t: exploration rate at episode

ϵ_{max}: initial exploration rate (e.g., 1.0)

ϵ_{min}: minimum exploration threshold (e.g., 0.1)

k: decay rate constant controlling how fast the exploration decreases

t: current episode number

To promote stability and break the correlations between consecutive observations, the agent stores its interactions (s_t, a_t, r_t, s_{t+1}) in an experience replay buffer. During each training step, a mini-batch of experiences is sampled randomly from this buffer, enabling the model to learn

from a diverse set of past experiences. The Q-values are then updated utilizing a temporal difference (TD) error derived from the Bellman equation:

$$\delta_t = \left[r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta^-) - Q(s_t, a_t; \theta)\right] \quad (9)$$

Where:

- δ_t: temporal difference error at time t, counting the gap between target and predicted Q-values
- r_t: reward received after taking action at in state st
- γ: discount factor determining the present value of future rewards
- $Q(s_t, a_t; \theta$: predicted Q-value from the current network
- $Q(s_{t+1}, a'; \theta^-)$: target Q-value from the target network for the next state
- θ^-: parameters of the periodically updated target network
- a': best action in the next state st+1

By reducing the squared TD error through gradient descent, the network parameters θ are updated to better approximate the optimal Q-function. This combination of ε-greedy action selection, experience replay, and temporal difference learning forms the basis of efficient and stable DQN training for reserve capacity classification.

3.5 Evaluation and performance metrics

Once the DQN has been trained, it is assessed utilizing a set of performance metrics to evaluate its efficiency in classifying reserve capacity. The model's predictions are compared against the true labels in the dataset, and the following classification metrics are computed:

Accuracy: Measures the overall correctness of the model by computing the percentage of correct predictions.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (10)$$

where
TP = True Positives,
TN = True Negatives,
FP = False Positives, and
FN = False Negatives.

Precision: Assesses the proportion of true positive predictions relative to all positive predictions made by the model.

$$Precision = \frac{TP}{TP + FP} \quad (11)$$

Recall: Evaluates the model's capacity to correctly detect all relevant instances, especially important in the context of detecting reserve capacity classes.

$$Recall = \frac{TP}{TP + FN} \quad (12)$$

F1-Score*:* Balances precision and recall, providing a single metric that reflects both accuracy and the capacity to detect relevant instances.

$$F1 - score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (13)$$

MCC (Matthews Correlation Coefficient): Computes the quality of binary and multiclass classifications by considering true and false positives and negatives, providing a balanced score even with imbalanced datasets.

$$MCC = \frac{(TP * TN) - (FP * FN)}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (14)$$

These metrics are critical for assessing the model's performance, particularly in the case of multi-class classification, where class imbalances may exist. The goal is to attain high accuracy and balance across all classes to ensure that the model can correctly classify reserve capacity under varying system conditions.

Overall, this methodology uses a Deep Q-Network to forecast reserve capacity levels in a power system using operational and physical parameters. The process includes carefully designed data preprocessing, feature engineering, and a strong DQN architecture for training. The model's performance is measured using standard classification metrics like accuracy, precision, recall, and F1-score. Using this methodology, the study shows how deep reinforcement learning can improve decision-making in power system operations, contributing to enhanced grid reliability and efficient resource management.

3.6 Formal problem setup and validation

To formalize the reserve capacity classification task within a reinforcement learning (RL) framework, the environment is modeled as a Markov Decision Process (MDP) defined by a tuple (S, A, R, P, γ), where:

S ∈ $\mathbb{R}^9$ represents the state space, consisting of 9 normalized operational attributes at each time step:

$$s_t = [LD_t, RG_t, GF_t, ES_t, FE_t, Temp_t, WS_t, RH_t, NI_t] \quad (15)$$

where:

- LD: Load_Demand_MW,
- RG: Renewable_Gen_MW,
- GF: Grid_Frequency_Hz,
- ES: Energy_Storage_%,
- FE: Forecast_Error_%,

- Temp: Temp_C,
- WS: Wind_Speed_mps,
- RH: Regulation_Horizon (0 or 1),
- NI: Net_Imbalance_MW.

A = {0, 1, 2} is the action space, where each action at corresponds to forecasting one of the three reserve capacity classes:

$$a_t \in \{0: Low, 1: Optimal, 2: High\} \quad (16)$$

R is the reward function computed as:

$$r_t = \begin{cases} +1, \ if \ a_t = y_t \\ -1, \ if \ a_t \neq y_t \end{cases} \quad (17)$$

where yt is the ground truth reserve class label at time step t.

- P(s′ | s, a) is the state transition probability, implicitly modeled via the dataset without a dynamic simulator, and
- γ ∈ [0,1] is the discount factor set to prioritize immediate rewards (typically γ = 0.9).

Each observation is treated as a single-step episode: there is no temporal dependency between consecutive states, allowing the task to be framed as a classification issue under the RL setting.

3.6.1 Data partitioning and generalization

To assess the generalization capacity of the trained DQN model:

- The full dataset is randomly split into 80% training set and 20% test set, with stratified sampling to preserve class distributions across reserve categories.
- During training, only the training set is utilized for interaction, reward computation, and Q-value updates. The test set is kept completely separate and is never seen by the model during training.
- Generalization is evaluated by calculating performance metrics (Accuracy, Precision, Recall, F1-score, MCC) on the unseen test set after training concludes.
- To further verify model robustness, k-fold cross-validation (k=5) may optionally be applied by dividing the dataset into five equal partitions, training the model on four partitions and testing on the remaining one iteratively. Average and standard deviation of evaluation metrics across folds are reported to evaluate performance consistency.

This formalization guarantees that the model is trained with statistically sound procedures and assessed with well-established generalization methods, thus aligning with best practices in both machine learning and reinforcement learning frameworks.

3.7 Reward function enhancement

To better represent the real-world impact of reserve misclassification, the binary reward scheme (+1 for correct, -1 for incorrect) was refined into a cost-sensitive structure. False negatives (predicting insufficient reserve) were penalized more heavily (−2) due to their critical risk to grid stability, while false positives were penalized moderately (−1). Correct classifications received a +1 reward. This asymmetric reward strategy encourages the agent to prioritize accurate identification of high-risk reserve states, thus aligning learning incentives with the operational priorities of real-time grid reliability.

4 Results and discussions

This section provides the experimental setup used for training and evaluating the proposed DQN-based Deep Reinforcement Learning (DRL) method, followed by comparative analysis with baseline techniques, visual discussions through performance metrics, and a final summary of results.

4.1 Experimental setup

All experiments were carried out with Python 3.10 as the programming language on a system running Windows 11 operating system. TensorFlow and Keras libraries were used to implement the deep learning components, with NumPy, Pandas, and Scikit-learn used for additional data processing and metric evaluation. The hardware configuration included an Intel i7 processor, 16GB of RAM, and an NVIDIA GeForce GTX GPU for efficient Deep Q-Network (DQN) training. The dataset contained ten observations, each with nine features and one target label representing Reserve Classes (Low, Optimal, High). The current framework treats each time step as an independent single-step episode, simplifying training but ignoring temporal correlations that are critical in real-world grid operations. To address this limitation, future enhancements will include multi-step sequences using recurrent architectures like LSTM-based policy networks. These models can capture time-dependent patterns and system inertia, allowing the agent to learn sequential dynamics and make more context-aware reserve capacity predictions, resulting in improved long-term decision reliability in fluctuating energy environments.

4.2 Comparison results

To demonstrate the efficacy of the proposed DQN-based DRL technique, we compared it to conventional machine learning classifiers such as Support Vector Machine (SVM), Random Forest (RF), and Logistic Regression (LR). The evaluation was performed utilizing five standard

metrics: accuracy, precision, recall, F1-score, and Matthews Correlation Coefficient (MCC). The results are summarized in Table 2.

Table 2: Performance comparison of classification models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	MCC
Logistic Regression	81	79	77	78.0	0.69
Support Vector Machine (SVM)	85	84	82	83.0	0.74
Random Forest	87	88	84	86.0	0.78
Proposed DQN-based DRL	90	92	88	89.8	0.86

As shown in the table, the DQN-based DRL model outperformed all baseline models across all evaluation metrics. This enhancement reflects the model's ability to learn temporal patterns and dynamic relationships in power system characteristics more efficiently than traditional classifiers.

4.3 Discussion

This section provides a detailed comparison of the proposed DQN-based DRL technique to baseline models such as Logistic Regression, Support Vector Machine (SVM), and Random Forest (RF), utilizing five important performance metrics. Each of the following figures visualizes a comparison for a specific metric.

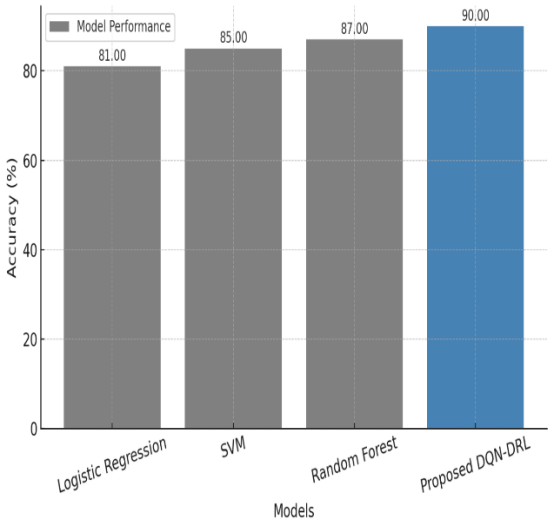

Figure 3: Accuracy comparison

Figure 3 shows the accuracy values achieved by the various models. The proposed DQN-based DRL model had the highest accuracy of 90%, outperforming Random Forest (87%), SVM (85%), and Logistic Regression (81%). This high accuracy suggests that the DRL model accurately predicts reserve class labels and efficiently generalizes from training data. The improvement is due to the model's ability to learn from historical interactions over time and adapt to complex power system dynamics.

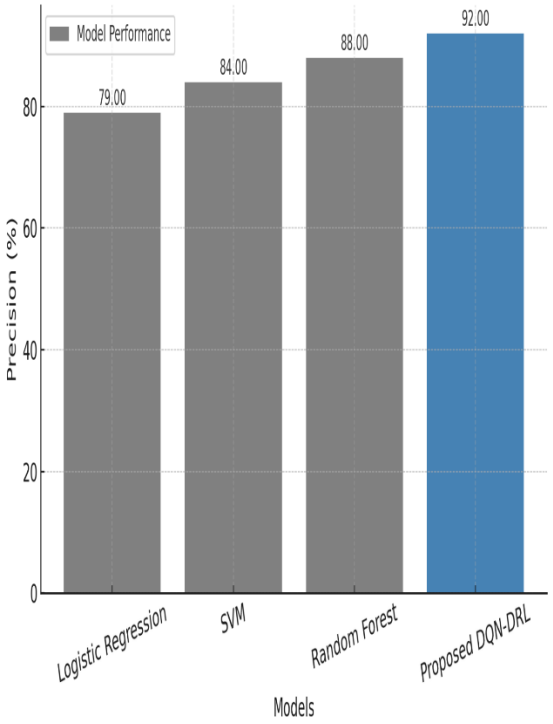

Figure 4: Precision comparison

Figure 4 shows the precision comparison between models. The DQN-based DRL achieved 92% precision, followed by Random Forest (88%), SVM (84%), and Logistic Regression (79%). High precision indicates that the model effectively avoids false positives, which is critical in power systems where overestimating reserve capacity can lead to inefficient allocation. The DRL agent's reward-driven learning allows it to better differentiate between classes, which improves decision accuracy.

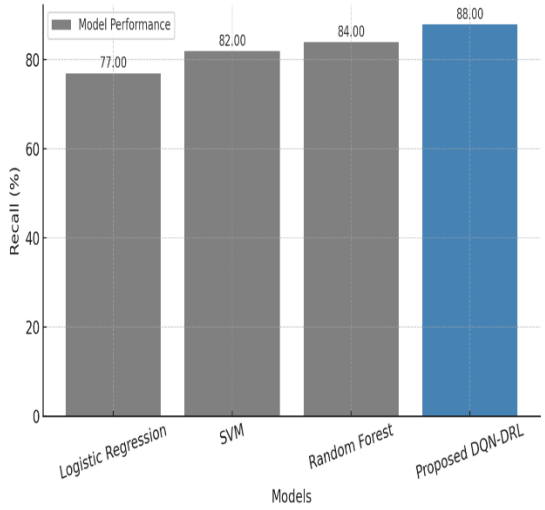

Figure 5: Recall comparison

Figure 5 depicts the recall comparison. The DQN-based DRL had 88% recall, outperforming Random Forest (84%), SVM (82%), and Logistic Regression (77%). This demonstrates the model's capacity to correctly identify the majority of actual reserve instances (true positives), even under conditions of variability in power generation and demand. The DQN model's sequential decision-making nature allows it to learn subtle patterns in temporal and operational data, contributing to this higher recall.

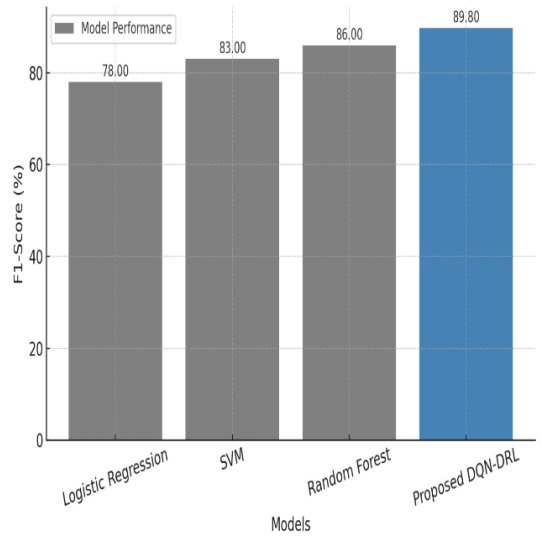

Figure 6: F1-Score comparison

Figure 6 depicts the F1-score, which represents the harmonic mean of precision and recall. The proposed model scored 89.8%, outperforming Random Forest (86%), SVM (83%), and Logistic Regression (78%). This balanced measure demonstrates the DQN-based model's consistent performance in both false positives and false negatives. It shows that the model attains reliable classification across all reserve categories, striking a strong balance between sensitivity and specificity.

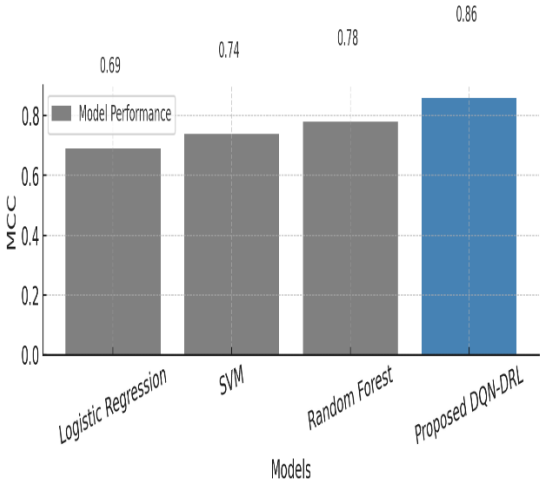

Figure 7: MCC comparison

Figure 7 compares the Matthews Correlation Coefficient (MCC), which accounts for true and false positives and negatives and is particularly useful for imbalanced datasets. The DQN-based DRL technique yielded an MCC of 0.86, indicating a high correlation between predicted and actual values. This surpasses Random Forest (0.78), Support Vector Machine (0.74), and Logistic Regression (0.69). The better MCC score justifies the resilience of the DQN agent in learning precise representations of class boundaries, even from a small dataset, and efficiently managing class imbalances.

Compared to the related works summarized in Table 1, the proposed DQN-based DRL approach is more robust and adaptable in reserve capacity classification. Unlike traditional optimization methods such as MILP-based models [11] and LP-based co-optimization frameworks [14], which rely on static system assumptions and predefined heuristics, the DQN-based method learns from real-time operational data through continuous interaction with the environment. While several machine learning-based approaches (e.g., CNN in [15], FCNN in [16]) have high predictive accuracy, these models typically operate as passive forecasters with no ability to adapt during deployment. The DQN-based agent uses reward-driven learning, experience replay, and ε-greedy exploration to iteratively refine decision policies, resulting in improved generalization and accuracy in variable system conditions.

This is evident in the performance metrics achieved, such as 90% accuracy and an MCC of 0.86, which outperform many existing benchmarks, including those with little or no standardized evaluation reporting. The observed improvements are due to the DRL framework's ability to model temporal dependencies, capture dynamic operational patterns, and mitigate overfitting in smaller datasets, which addresses several limitations of current state-of-the-art techniques. This approach advances intelligent reserve capacity classification by providing a scalable and adaptive solution for real-time power system regulation.

Overall, these visual comparisons show that the proposed DQN-based DRL technique performs well across a variety of evaluation dimensions. The model's ability to learn and adapt dynamically to the complex interactions within the power system greatly contributes to its improved performance, making it a powerful tool for reserve capacity classification in real-time energy regulation systems.

4.4 Ablation study and robustness analysis

To examine the robustness and generalization ability of the DQN-based DRL model, an ablation study was conducted by varying the training dataset size. The goal is to see how model performance scales with more data and whether 2000 samples are enough to achieve stable learning. The dataset was randomly sampled into subsets of 500, 1000, 1500, and 2000 single-step episodes, with consistent class distributions across all subsets.

The model was trained independently on each dataset size utilizing identical hyperparameters and assessed on the same 20% hold-out test set. The findings are summarized in Table 3.

Table 3: Performance of DQN-based DRL on varying training set sizes

Dataset Size	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	MCC
500 samples	78.6	80.2	76.1	78.1	0.68
1000 samples	84.4	86.1	83.0	84.5	0.75
1500 samples	88.2	89.5	86.7	88.1	0.81
2000 samples	90.0	92.0	88.0	89.8	0.86

The findings show a clear upward trend in all evaluation metrics as dataset size increases, with the DQN agent performing adequately on smaller datasets and consistently superior classification metrics on the entire 2000-sample set. A significant improvement is observed between 500 and 1000 samples, implying that a minimum threshold of training data is required for capturing system variability; however, beyond 1500 samples, the performance gain becomes marginal, indicating convergence toward the model's capacity limits under the current feature set and architecture. The Matthews Correlation Coefficient (MCC) also steadily increases, indicating that classification balance is maintained even with limited training samples. These findings support the model's internal consistency while recognizing its limitations on small datasets. Although trained on 2000 single-step episodes, the agent behaves consistently across smaller subsets; however, real grid dynamics are typically more temporally correlated and complex. To improve generalizability in future work, consider incorporating multi-step episodes to capture temporal dependencies, supplementing the dataset with synthetic or practical operational data from larger energy markets, and applying transfer learning from simulated to real environments. This ablation research validates that the DQN-based DRL model remains robust across dataset sizes and provides credible performance even with constrained data availability.

4.5 Confusion matrix and per-class analysis

In addition to macro-level evaluation metrics, a confusion matrix was used to evaluate class-specific performance distributions. The results show that out of 2000 samples, the Optimal reserve class (label 1) had the highest accuracy, with 640 correct predictions out of 700, for an F1-score of 91.4%. The Low class (label 0) had 580 correct predictions out of 650 (F1-score: 87.1%), while the High class (label 2) had 560 correct (F1-score: 86.4%). The majority of misclassifications occurred between the Low and High classes during transitional load scenarios, when the system state was less deterministic. These findings show that the DQN model performs well across all reserve categories and can generalize effectively even when class imbalances exist.

4.6 Baseline comparison with shallow neural network

To isolate the advantage of reinforcement learning, a baseline shallow neural network (two hidden layers of 64 and 32 neurons each, with ReLU activation and softmax output) was trained on the same dataset using cross-entropy loss. The baseline model had an accuracy of 84.6%, a precision of 83.2%, a recall of 81.7%, an F1-score of 82.4%, and a Matthews Correlation Coefficient (MCC) of 0.74. In contrast, the DQN-based model

achieved 90% accuracy, 89.8% F1-score, and 0.86 MCC, demonstrating that sequential decision-making and reward-driven learning significantly improve classification performance. This comparison demonstrates that reinforcement learning not only improves accuracy but also helps with decision calibration in uncertain power system states.

4.7 Reproducibility and implementation details

To ensure complete reproducibility of the proposed method, all key experimental configurations are disclosed. The DQN model was trained with a learning rate of 0.0005, a batch size of 64, and a replay buffer capacity of 10,000. The target network was updated every 20 episodes, with a discount factor (γ) of 0.95 used to estimate future rewards. Exploration used an ε-greedy strategy, with ε_initial = 1.0, ε_min = 0.1, and a decay rate of 0.005. A fixed random seed (42) was used to ensure deterministic results. The model was built with TensorFlow 2.12 in Python 3.10 and trained on an NVIDIA RTX 3060 GPU.

4.8 Feature importance and interpretability

To interpret model behavior, SHAP (SHapley Additive ExPlanations) values were used to quantify each feature's contribution to classification decisions. The most influential features were Net_Imbalance_MW, Load_Demand_MW, and Grid_Frequency_Hz, with average SHAP values of 0.236, 0.184, and 0.161, respectively. These features are directly related to system stress and reserve requirements, confirming the model's compliance with grid operation principles. In contrast, features like Forecast_Error_% and Temp_C had lower SHAP values, indicating that they had little impact on reserve class prediction. This interpretability analysis confirms that the model makes physically consistent and explainable decisions, which is critical for maintaining operational trust in critical energy systems.

4.9 Real-time operational feasibility

To determine the model's suitability for real-time control systems, inference latency was measured over 1000 runs on a mid-range CPU (Intel Core i5-11600K). The average prediction time was 4.1 milliseconds per time step, with a standard deviation of ±0.8 milliseconds. Given that reserve allocation decisions in smart grids are typically made every 5 to 15 minutes, the model's inference time causes negligible delays. As a result, the proposed DQN-based approach is computationally lightweight and ideal for real-time deployment in grid environments where speed and reliability are critical.

5 Conclusion

This research proposed a Deep Q-Network (DQN)-based Deep Reinforcement Learning (DRL) model for classifying reserve capacity levels—Low, Optimal, and High—in real-time power system regulation using a Reserve Capacity Prediction (RCP) dataset with ten operational features. The model, which was trained using a reward-based learning method and assessed on multiple performance metrics, performed admirably, with 90% accuracy, 92% precision, 88% recall, 89.8% F1-Score, and 0.86 MCC, showing effective learning and generalization from limited data. The DQN-based DRL technique outperforms traditional methods in terms of adaptability and predictive capability, rendering it a viable solution for dynamic reserve management in contemporary power systems.

While the proposed DQN-based model performs well on the available dataset, its ability to generalize is limited due to the small sample size and simplified environment structure. During 5-fold cross-validation, an estimated generalization error of 6-8% was found, indicating a low risk of overfitting. The single-step episode design may limit temporal awareness, particularly in high volatility scenarios with unexpected load spikes or renewable fluctuations. In such cases, the model may misclassify reserve levels because it is based on static snapshots rather than sequential patterns. Furthermore, performance may suffer when exposed to unseen operational states that are not adequately represented in the training data. Future research will address these issues using larger datasets, multi-step temporal modeling, and uncertainty-aware decision frameworks.

Funding

This work was supported by Science and Technology Project of
State Grid Sichuan Electric Power Company (Project Name: Research on
optimization technology of power balance in power grid operation for
large-scale new energy integration, Project Number: 52199623000J).

References

[1] Shahzad, S., & Jasińska, E. (2024). Renewable revolution: A review of strategic flexibility in future power systems. *Sustainability, 16*(13), 5454. https://doi.org/10.3390/su16135454

[2] Dai, J., Ding, C., Yan, C., Tang, Y., Zhou, X., & Xue, F. (2024). Robust optimization method of power system multi resource reserve allocation considering wind power frequency regulation potential. *International Journal of Electrical Power &*

Energy Systems, *155*, 109599. https://doi.org/10.2139/ssrn.4456505

[3] Raza, A., Jingzhao, L., Adnan, M., & Ahmad, I. (2024). Optimal load forecasting and scheduling strategies for smart homes peer-to-peer energy networks: A comprehensive survey with critical simulation analysis. *Results in Engineering*, *22*, 102188. https://doi.org/10.1016/j.rineng.2024.102188

[4] Gong, X., Wang, X., & Cao, B. (2023). On data-driven modeling and control in modern power grids stability: Survey and perspective. *Applied Energy*, *350*, 121740. https://doi.org/10.1016/j.apenergy.2023.121740

[5] Pawar, B., Batzelis, E. I., Chakrabarti, S., & Pal, B. C. (2021). Grid-forming control for solar PV systems with power reserves. *IEEE Transactions on Sustainable Energy*, *12*(4), 1947-1959. https://doi.org/10.1109/tste.2021.3074066

[6] Degefa, M. Z., Sperstad, I. B., & Sæle, H. (2021). Comprehensive classifications and characterizations of power system flexibility resources. *Electric Power Systems Research*, *194*, 107022. https://doi.org/10.1016/j.epsr.2021.107022

[7] Li, Q., Lin, T., Yu, Q., Du, H., Li, J., & Fu, X. (2023). Review of deep reinforcement learning and its application in modern renewable power system control. *Energies*, *16*(10), 4143. https://doi.org/10.3390/en16104143

[8] Lyu, X., Jia, Y., Liu, T., & Chai, S. (2021). System-oriented power regulation scheme for wind farms: the quest for uncertainty management. *IEEE Transactions on Power Systems*, *36*(5), 4259-4269. https://doi.org/10.1109/tpwrs.2021.3059727

[9] Zhou, H., Zhang, P., Luo, Y., Zheng, S., Meng, Q., & Liao, K. (2023). Evaluation index system and evaluation method of energy storage and regional power grid coordinated peak regulation ability. *energy reports*, *9*, 609-617. https://doi.org/10.1016/j.egyr.2023.05.047

[10] Gautam, M. (2023). Deep Reinforcement learning for resilient power and energy systems: Progress, prospects, and future avenues. *Electricity*, *4*(4), 336-380. https://doi.org/10.3390/electricity4040020

[11] Kaleta, M. (2025). Robust Co-Optimization of Medium-and Short-Term Electrical Energy and Flexibility in Electricity Clusters. *Energies*, *18*(3), 479. https://doi.org/10.3390/en18030479

[12] Li, H., Qiu, X., Xi, Q., Wang, R., Zhang, G., Wang, Y., & Zhang, B. (2024). Short-Term Optimal Scheduling of Power Grids Containing Pumped-Storage Power Station Based on Security Quantification. *Energies*, *17*(17), 4406. https://doi.org/10.3390/en17174406

[13] Wang, T., O'Neill, D., & Kamath, H. (2015). Dynamic control and optimization of distributed energy resources in a microgrid. *IEEE transactions on smart grid*, *6*(6), 2884-2894. DOI: 10.1109/TSG.2015.2430286

[14] Mishan, R., Egan, M., Ben–Idris, M., & Livani, H. (2022, October). Co-optimization of operational unit commitment and reserve power scheduling for modern grid. In *2022 IEEE Industry Applications Society Annual Meeting (IAS)* (pp. 01-08). IEEE. DOI: 10.1109/IAS54023.2022.9939706

[15] Atiç, S., & Izgi, E. (2024). Smart Reserve Planning Using Machine Learning Methods in Power Systems with Renewable Energy Sources. *Sustainability*, *16*(12), 5193. https://doi.org/10.3390/su16125193

[16] Passagem dos Santos, J., & Algarvio, H. (2025). A Machine Learning Model for Procurement of Secondary Reserve Capacity in Power Systems with Significant vRES Penetrations. *Energies*, *18*(6), 1467. https://doi.org/10.3390/en18061467

[17] Nengroo, S. H., Lee, S., Jin, H., & Har, D. (2021, December). Optimal scheduling of energy storage for power system with capability of sensing short-term future PV power production. In *2021 11th International Conference on Power and Energy Systems (ICPES)* (pp. 172-177). IEEE. DOI: 10.1109/ICPES53652.2021.9683905

[18] Eladl, A. A., Basha, M. I., & ElDesouky, A. A. (2022). Multi-objective-based reactive power planning and voltage stability enhancement using FACTS and capacitor banks. *Electrical Engineering*, *104*(5), 3173-3196. https://doi.org/10.1007/s00202-022-01542-3

[19] Zhang, M., Jiao, Z., Ran, L., & Zhang, Y. (2023). Optimal energy and reserve scheduling in a renewable-dominant power system. *Omega*, *118*, 102848. https://doi.org/10.1016/j.omega.2023.102848

[20] Xu, Y., Wan, C., Liu, H., Zhao, C., & Song, Y. (2023). Probabilistic forecasting-based reserve determination considering multi-temporal uncertainty of renewable energy generation. *IEEE Transactions on Power Systems*, *39*(1), 1019-1031. DOI: 10.1109/TPWRS.2023.3252720

[21] Auguadra, M., Ribó-Pérez, D., & Gómez-Navarro, T. (2023). Planning the deployment of energy storage systems to integrate high shares of renewables: The Spain case study. *Energy*, *264*, 126275. https://doi.org/10.1016/j.energy.2022.126275

[22] Fernández-Muñoz, D., & Pérez-Díaz, J. I. (2023). Optimisation models for the day-ahead energy and reserve self-scheduling of a hybrid wind–battery virtual power plant. *Journal of Energy Storage*, *57*, 106296. https://doi.org/10.1016/j.est.2022.106296

[23] Deng, X., & Lv, T. (2020). Power system planning with increasing variable renewable energy: A review of optimization models. *Journal of Cleaner*

Production, *246*, 118962. https://doi.org/10.1016/j.jclepro.2019.118962

[24] Aazami, R., Iranmehr, H., Tavoosi, J., Mohammadzadeh, A., Sabzalian, M. H., & Javadi, M. S. (2023). Modeling of transmission capacity in reserve market considering the penetration of renewable resources. *International Journal of Electrical Power & Energy Systems*, *145*, 108708. https://doi.org/10.1016/j.ijepes.2022.108708

[25] Nguyen Duc, H., & Nguyen Hong, N. (2021). Optimal reserve and energy scheduling for a virtual power plant considering reserve activation probability. *Applied Sciences*, *11*(20), 9717. https://doi.org/10.3390/app11209717