

OFFLINE REINFORCEMENT LEARNING FOR MICROGRID VOLTAGE REGULATION

Shan Yang

School of System Science and Engineering
Sun-Yat Sen University
Guangzhou, China
yangsh237@mail2.sysu.edu.cn

Yongli Zhu*

School of System Science and Engineering
Sun-Yat Sen University
Guangzhou, China
yzhu16@alum.utk.edu

ABSTRACT

This paper presents a study on using different offline reinforcement learning algorithms for microgrid voltage regulation with solar power penetration. When environment interaction is unviable due to technical or safety reasons, the proposed approach can still obtain an applicable model through offline-style training on a previously collected dataset, lowering the negative impact of lacking online environment interactions. Experiment results on the IEEE 33-bus system demonstrate the feasibility and effectiveness of the proposed approach on different offline datasets, including the one with merely low-quality experience.

1 INTRODUCTION

The rapid integration of distributed energy resources (DERs), such as photovoltaic (PV) systems, has introduced significant challenges in maintaining stable voltage levels within power distribution networks (Tonkoski et al., 2012). Unlike traditional centralized power systems, modern distribution networks, e.g., microgrids with multiple distributed renewable sources, are more dynamic and decentralized Ali et al. (2024); Fusco and Russo (2021), making voltage regulation particularly challenging. Advanced control strategies are required to address fluctuations caused by changes in load demand and intermittent renewable generation (Wei et al., 2024; Singhal et al., 2018).

Reinforcement learning (RL) has emerged as a powerful tool for optimizing control policies in complex and dynamic environments (Cao et al., 2020). Its ability to learn from past experiences and adapt to diverse scenarios makes it an appealing solution for voltage regulation. However, standard RL methods, such as Q-learning and Deep Deterministic Policy Gradient (DDPG), typically rely on real-time environmental interactions to explore and improve their policies. While effective, this reliance on real-time interactions makes these methods challenging to apply in real-world power systems, where online experimentation can be costly, risky, and potentially destabilizing (Levine et al., 2020; Schmidt et al., 2024).

Offline reinforcement learning (Offline RL) trains agents on pre-collected datasets, reducing the need for real-time interactions (Chen et al., 2024; Hu et al., 2024). However, it faces challenges like extrapolation errors when evaluating out-of-distribution actions, leading to overly optimistic Q-values and unstable policies (Levine et al., 2020). These risks are critical in safety-sensitive tasks like voltage regulation (Zhang et al., 2024). To address this, algorithms such as batch-constrained deep Q-learning (BCQ) constrain actions to remain close to the offline dataset, reducing the likelihood of high-risk decisions (Fujimoto et al., 2019).

This study compares the BCQ and CQL (Kumar et al., 2020) algorithms in addressing the voltage regulation problem in a PV-penetrated distribution network. Section 2 details the architectures of both the BCQ and CQL algorithmic frameworks, highlighting their respective mechanisms for mitigating extrapolation errors. Section 3 describes the offline dataset creation and collection used in this study, including datasets of varying quality levels. Section 4 outlines the experimental setup and presents simulation results on the IEEE 33-bus system, focusing on the performance of BCQ

*Corresponding author.

and CQL across different datasets. Finally, Section 5 provides conclusions and discusses potential future work in this research area.

2 OFFLINE REINFORCEMENT LEARNING

2.1 BASIC IDEA

In scenarios such as voltage regulation of a microgrid, direct experimentation with the live system is often impractical due to cost, safety, and legal concerns.

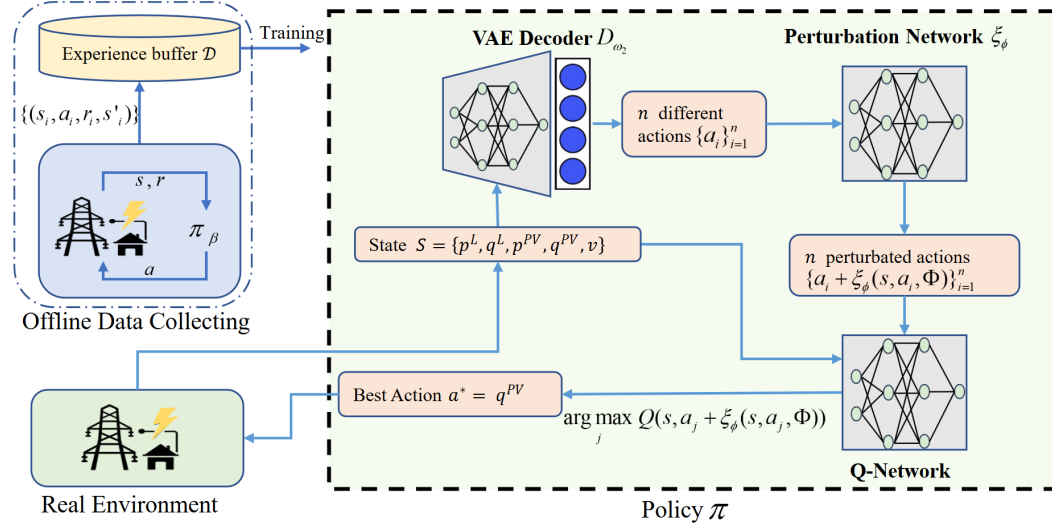


Figure 1: The basic idea of offline reinforcement learning.

Offline reinforcement learning addresses these challenges by training solely on pre-collected datasets ($\mathcal{D} = \{(s^i, a^i, r^i, s_{i+1}^i)\}$, where transitions are gathered by a behavior policy π_β). This data-driven approach is suitable for settings where system safety and stability are critical and typically involves three steps:

- (i) Data collection using a fixed (or simple) behavior policy;
- (ii) Offline training with the collected dataset;
- (iii) Deployment of the trained policy in the real environment.

This approach allows the agent to develop a policy that meets control goals while avoiding the risks of direct interaction with the system.

2.2 THE BCQ ALGORITHM

The BCQ algorithm intends to maximize the expected cumulative reward under a constraint that prevents the agent from selecting actions significantly different from those observed in the offline dataset. In an offline setting, this is crucial for preventing the agent from overestimating Q-values for unfamiliar state-action pairs. Given a dataset \mathcal{D} collected by a behavior policy π_β , the goal of BCQ is to learn a policy π that maximizes:

$$\mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} \left[r + \gamma \max_{a' \in \text{Batch}} Q(s', a') \right] \quad (1)$$

where γ is the discount factor, and “Batch” refers to the set of actions close to the distribution in \mathcal{D} .

To further constrain the action selection and improve generalization, BCQ uses a state-conditioned generative model, namely a variational autoencoder (VAE) (denoted by $G_\omega(s)$), to approximate feasible actions for each state. The VAE is trained to generate actions that are likely under the behavior policy π_β , filtering out actions that deviate too far from the observed data. This setup effectively restricts the agent’s exploration to actions close to those in the offline dataset, mitigating the risk of extrapolation errors. More algorithmic details can be found in Appendix A.1.1.

2.3 THE CQL ALGORITHM

Conservative Q-Learning (CQL) is an offline reinforcement learning algorithm designed to mitigate extrapolation errors by modifying the Q-function’s loss function. Extrapolation errors occur when the Q-function assigns overly optimistic values to actions that are poorly represented or absent in the dataset, leading to unreliable or unsafe policies. To address this, CQL introduces an additional penalty term to the Q-value loss, ensuring conservative estimates and prioritizing stable policy outcomes.

CQL optimizes the following objective by penalizing overestimation:

$$J_{\text{CQL}}(Q) = \mathbb{E}_{s \sim \mathcal{D}} \left[\log \sum_a \exp Q(s, a) - \mathbb{E}_{a \sim \mathcal{D}} [Q(s, a)] \right] \quad (2)$$

where the first term penalizes high Q-values across the entire action space, preventing overestimation for unobserved actions. The second term prioritizes actions within the dataset distribution to ensure conservative policy behavior. More algorithmic details can be found in Appendix A.1.2.

3 MICROGRID VOLTAGE REGULATION

3.1 ENVIRONMENT

The environment is a `distflow` program that simulates the power grid’s power flow at each time step (Wang et al., 2021). It evaluates control algorithms under realistic operational conditions by simulating variations in load and PV generation. The environment includes state, action, and reward designs tailored for the voltage regulation task. For brevity, the detailed definitions of the state space, action space, reward function, and objective function are provided in Appendix A.2.

3.2 DATA COLLECTED BY DIFFERENT LEVELS OF POLICIES

We create three distinct datasets: *Expert*, *Medium*, and *Poor*. Each dataset contains 200000 transitions (s, a, r, s') .



Figure 2: The episode returns across datasets.

Expert Dataset This dataset is generated by a well-trained DDPG agent that consistently performs at a near-optimal level. The Expert dataset contains high-quality transitions with high average returns, simulating the scenario where the agent has a “good teacher”. This dataset benchmarks how well an agent can perform when provided with the best possible data.

Medium Dataset Created by introducing Gaussian noise (5%) to the Expert dataset’s actions, this dataset simulates a moderately effective policy. The added noise introduces variability, resulting in a mix of high and moderate-performance episodes. This dataset broadens the state-action pairs the agent encounters, testing its ability to learn from data that include suboptimal experiences.

Poor Dataset This dataset *deliberately* employs a poor policy by randomly selecting actions from the action space. It includes episodes of low average returns due to the lack of purposeful actions. This dataset pushes the offline-RL agent to extract useful patterns from poorly performed data as much as possible. It is used as a touchstone to inspect the genuine effectiveness of the offline RL algorithm before it is applied to the real system.

4 EXPERIMENT RESULTS

This section presents the experimental results of respectively applying the BCQ and CQL algorithms for the voltage regulation task on the IEEE 33-bus system (cf. Appendix A.5).

This section compares BCQ and CQL on three offline datasets: Poor, Medium, and Expert. The results, shown in Fig. 3, highlight each algorithm’s performance in maintaining controllable voltage levels and minimizing voltage deviations. For more details on the performance of BCQ across the datasets, please refer to Appendix A.6.

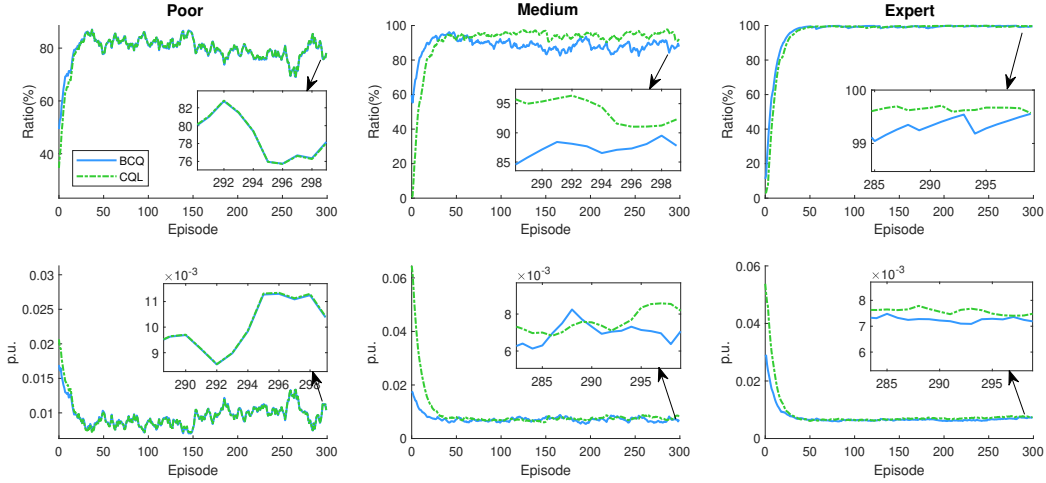


Figure 3: Performance comparison of BCQ and CQL across three datasets: Totally Controllable Ratio (top row) and Average Voltage Deviation (bottom row).

The comparison focuses on two key metrics: 1) *Totally Controllable Ratio*: The percentage of time steps where all buses maintain voltages within the acceptable range (0.95 to 1.05 p.u.), reflecting controllability and stability. 2) *Average Voltage Deviation*: The average deviation of bus voltages from the target level (1.0 p.u.), indicating voltage regulation effectiveness.

Performance on Poor Dataset: CQL achieves higher controllable ratios and lower voltage deviations than BCQ, demonstrating robustness on low-quality data

Performance on Medium Dataset: CQL significantly outperforms BCQ, benefiting from its emphasis on the in-distribution actions, leading to better stability even with moderate data noise.

Performance on Expert Dataset: Both algorithms perform well with high-quality data, achieving high controllable ratios and low deviations. CQL behaves slightly more stable.

In summary, CQL outperforms BCQ across all three datasets, particularly in scenarios with low and moderate data quality, making it a safer and more robust choice for voltage regulation tasks.

5 CONCLUSION

This study investigates the application of Offline Reinforcement Learning (Offline RL) for voltage regulation in the PV-penetrated microgrid, focusing on BCQ and CQL algorithms. The results demonstrate that Offline RL can effectively utilize pre-collected datasets to learn stable and reliable control policies without requiring real-time interactions.

CQL consistently outperforms BCQ across datasets of varying quality, particularly in scenarios with low or medium-quality data, due to its conservative Q-value estimation that enhances robustness and stability. While BCQ provides a solid baseline, its performance is more sensitive to data quality.

Future research could explore applying these methods to larger and more complex power grids, integrating additional Offline RL approaches, and improving scalability to address real-world operational challenges in power systems.

REFERENCES

- Arman Ali, Chaojie Li, and Branislav Hredzak. Dynamic voltage regulation in active distribution networks using day-ahead multi-agent deep reinforcement learning. *IEEE Transactions on Power Delivery*, 2024.
- Di Cao, Weihao Hu, Junbo Zhao, Guozhou Zhang, Bin Zhang, Zhou Liu, Zhe Chen, and Frede Blaabjerg. Reinforcement learning and its applications in modern power and energy systems: A review. *Journal of modern power systems and clean energy*, 8(6):1029–1042, 2020.
- Xiacong Chen, Siyu Wang, Julian McAuley, Dietmar Jannach, and Lina Yao. On the opportunities and challenges of offline reinforcement learning for recommender systems. *ACM Transactions on Information Systems*, 42(6):1–26, 2024.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, pages 2052–2062. PMLR, 2019.
- Giuseppe Fusco and Mario Russo. A decentralized approach for voltage control by multiple distributed energy resources. *IEEE Transactions on Smart Grid*, 12(4):3115–3127, 2021.
- Daner Hu, Zichen Li, Zhenhui Ye, Yonggang Peng, Wei Xi, and Tiantian Cai. Multi-agent graph reinforcement learning for decentralized volt-var control in power distribution systems. *International Journal of Electrical Power & Energy Systems*, 155:109531, 2024.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Carolin Schmidt, Daniele Gammelli, Francisco Camara Pereira, and Filipe Rodrigues. Learning to control autonomous fleets from observation via offline reinforcement learning. In *2024 European Control Conference (ECC)*, pages 1399–1406. IEEE, 2024.
- Ankit Singhal, Venkataramana Ajjarapu, Jason Fuller, and Jacob Hansen. Real-time local volt/var control under external disturbances with high pv penetration. *IEEE Transactions on Smart Grid*, 10(4):3849–3859, 2018.
- Reinaldo Tonkoski, Dave Turcotte, and Tarek HM El-Fouly. Impact of high pv penetration on voltage profiles in residential neighborhoods. *IEEE Transactions on Sustainable Energy*, 3(3): 518–527, 2012.
- Jianhong Wang, Wangkun Xu, Yunjie Gu, Wenbin Song, and Tim C Green. Multi-agent reinforcement learning for active voltage control on power distribution networks. *Advances in Neural Information Processing Systems*, 34:3271–3284, 2021.
- Xiang Wei, Xian Zhang, Guibin Wang, Ze Hu, Ziqing Zhu, and Ka Wing Chan. Online voltage control strategy: Multi-mode based data-driven approach for active distribution networks. *IEEE Transactions on Industry Applications*, 2024.
- Xiao Zhang, Zhi Wu, Qirun Sun, Wei Gu, Shu Zheng, and Jingtao Zhao. Application and progress of artificial intelligence technology in the field of distribution network voltage control: A review. *Renewable and Sustainable Energy Reviews*, 192:114282, 2024.

A APPENDIX

A.1 ALGORITHM DETAILS OF BCQ AND CQL

A.1.1 THE BCQ ALGORITHM

In the BCQ framework, the VAE outputs a set of candidate actions $\{a_i \sim G_\omega(s')\}_{i=1}^n$ for each state s' . These actions are then slightly perturbed to allow exploration within a safe neighborhood of the observed actions:

$$a_i = a_i + \xi_\phi(s', a_i, \Phi), \quad (3)$$

where ξ_ϕ is a perturbation model with a bound value Φ .

Finally, the Q-function in BCQ is updated by taking a conservative approach to avoid overestimation. The target value for a transition (s, a, r, s') is calculated as:

$$y = r + \gamma \max_{a_i \in \text{Batch}} \left[\lambda \min_{j=1,2} Q_{\theta_j}(s', a_i) + (1 - \lambda) \max_{j=1,2} Q_{\theta_j}(s', a_i) \right], \quad (4)$$

where λ controls the balance between pessimistic and optimistic value estimates, and Q_{θ_j} denotes the target Q-network.

By integrating the VAE to generate action candidates close to the dataset’s behavior, BCQ ensures a conservative learning approach, enhancing its suitability for the offline voltage regulation task in power grid. The complete pseudocode for the BCQ algorithm is shown in Algorithm 1.

Algorithm 1 Batch-Constrained Q-learning (BCQ)

Input: Batch B , horizon T , target update rate τ , mini-batch size N , max perturbation Φ , sampled actions n , weighting λ

Output: Trained networks for BCQ

- 1: Initialize $Q_{\theta_1}, Q_{\theta_2}, \xi_\phi, G_\omega = \{E_{\omega_1}, D_{\omega_2}\}$ with random parameters $\theta_1, \theta_2, \phi, \omega$
 - 2: Initialize target networks $Q_{\theta'_1}, Q_{\theta'_2}, \xi'_\phi: \theta'_1 \leftarrow \theta_1, \theta'_2 \leftarrow \theta_2, \phi' \leftarrow \phi$
 - 3: **for** $t = 1$ to T **do**
 - 4: Sample N transitions (s, a, r, s') from B and compute $\mu, \sigma = E_{\omega_1}(s, a), a' = D_{\omega_2}(s, z), z \sim \mathcal{N}(\mu, \sigma)$. Update $\omega \leftarrow \arg \min_\omega \sum (a - a')^2 + D_{\text{KL}}(\mathcal{N}(\mu, \sigma) \| \mathcal{N}(0, 1))$.
 - 5: Sample n actions $\{a_i \sim G_\omega(s')\}_{i=1}^n$, perturb $a_i: a_i \leftarrow a_i + \xi_\phi(s', a_i, \Phi)$.
 - 6: Compute target $y = r + \gamma \max_{a_i} \left[(1 - \lambda) \max_{j=1,2} Q_{\theta'_j}(s', a_i) + \lambda \min_{j=1,2} Q_{\theta'_j}(s', a_i) \right]$.
 - 7: Update $\theta_i \leftarrow \arg \min_{\theta_i} \sum (y - Q_{\theta_i}(s, a))^2$.
 - 8: Update $\phi \leftarrow \arg \max_\phi \sum Q_{\theta_i}(s, a + \xi_\phi(s, a, \Phi)), a \sim G_\omega(s)$.
 - 9: Update target networks: $\theta'_i \leftarrow \tau \theta_i + (1 - \tau) \theta'_i, \phi' \leftarrow \tau \phi + (1 - \tau) \phi'$.
 - 10: **end for**
-

A.1.2 THE CQL ALGORITHM

In the CQL framework, the following penalty is incorporated into the Q-function update (along with the standard Bellman loss) to make the training process more stable:

$$L(Q) = J_{\text{CQL}}(Q) + \zeta \cdot \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} \left[\left(r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right)^2 \right] \quad (5)$$

where ζ balances the conservative penalty $J_{\text{CQL}}(Q)$ and the Bellman consistency.

Furthermore, to encourage conservative behavior, CQL introduces an action-specific penalty term during policy evaluation, ensuring that the learned policy does not favor unobserved actions:

$$\pi(a | s) \approx \exp \left(Q(s, a) - \max_{a'} Q(s, a') \right) \quad (6)$$

where the softmax-like formulation adjusts action probabilities, giving higher weights to actions with lower penalties.

By combining these components, CQL effectively reduces the risk of unsafe policy execution, making it particularly suitable for offline reinforcement learning scenarios where live environment interaction is restricted. The complete pseudocode for the BCQ algorithm is shown in Algorithm 2.

Algorithm 2 Conservative Q-Learning (CQL)

Input: Batch B , horizon T , target update rate τ , mini-batch size N , max perturbation Φ , sampled actions n , weighting λ and ζ

Output: Trained networks for CQL

- 1: Initialize $Q_{\theta_1}, Q_{\theta_2}, \xi_\phi, G_\omega = \{E_{\omega_1}, D_{\omega_2}\}$ with random parameters $\theta_1, \theta_2, \phi, \omega$
- 2: Initialize target networks $Q_{\theta'_1}, Q_{\theta'_2}, \xi_{\phi'}: \theta'_1 \leftarrow \theta_1, \theta'_2 \leftarrow \theta_2, \phi' \leftarrow \phi$
- 3: **for** $t = 1$ to T **do**
- 4: Sample N transitions (s, a, r, s') from B and compute $\mu, \sigma = E_{\omega_1}(s, a), a' = D_{\omega_2}(s, z), z \sim \mathcal{N}(\mu, \sigma)$. Update $\omega \leftarrow \arg \min_\omega \sum (a - a')^2 + D_{\text{KL}}(\mathcal{N}(\mu, \sigma) \parallel \mathcal{N}(0, 1))$.
- 5: Sample n actions $\{a_i \sim G_\omega(s')\}_{i=1}^n$, perturb $a_i: a_i \leftarrow a_i + \xi_\phi(s', a_i, \Phi)$.
- 6: Compute target $y = r + \gamma \max_{a_i} \left[(1 - \lambda) \max_{j=1,2} Q_{\theta'_j}(s', a_i) + \lambda \min_{j=1,2} Q_{\theta'_j}(s', a_i) \right]$.
- 7: Compute CQL penalty:

$$J_{\text{CQL}}(Q_{\theta_i}) = \mathbb{E}_{s \sim B} \left[\log \sum_a \exp Q_{\theta_i}(s, a) - \mathbb{E}_{a \sim B} [Q_{\theta_i}(s, a)] \right]$$

- 8: Update Q-function with combined CQL penalty and Bellman loss:

$$L(Q_{\theta_i}) = J_{\text{CQL}}(Q_{\theta_i}) + \zeta \cdot \mathbb{E}_{(s,a,r,s') \sim B} \left[(y - Q_{\theta_i}(s, a))^2 \right]$$

- 9: Update $\theta_i \leftarrow \arg \min_{\theta_i} L(Q_{\theta_i})$.
- 10: Update $\phi \leftarrow \arg \max_\phi \sum Q_{\theta_i}(s, a + \xi_\phi(s, a, \Phi)), a \sim G_\omega(s)$.
- 11: Update target networks: $\theta'_i \leftarrow \tau \theta_i + (1 - \tau) \theta'_i, \phi' \leftarrow \tau \phi + (1 - \tau) \phi'$.
- 12: **end for**

A.2 ENVIRONMENT DETAILS

The environment is built on a `distflow` program that solves the microgrid’s power flow at each time step (Wang et al., 2021). By simulating realistic variations in load demand and PV generation, it provides a rigorous framework to evaluate the performance of control algorithms.

State Definition (S): The state S represents the current operational state of the microgrid and includes the following components:

- Active and reactive power loads at each bus: (p^L, q^L) ,
- Active and reactive power generated by PV inverters: (p^{PV}, q^{PV}) ,
- Voltage magnitudes at each bus: \mathbf{v} .

Mathematically, the state is defined as:

$$S = \{p^L, q^L, p^{PV}, q^{PV}, \mathbf{v}\} \quad (7)$$

Action Definition (A): The action a_k corresponds to the control signal for the k -th PV inverter, representing the ratio of its maximum reactive power output:

$$q_k^{PV} = a_k \cdot \sqrt{(S_k^{\text{max}})^2 - (p_k^{PV})^2} \quad (8)$$

where $-c \leq a_k \leq c$, and c is the maximum allowable ratio of reactive power injection or absorption. This design allows the agent to control voltage levels by adjusting the reactive power output of PV inverters.

Reward Function (r): The reward function incentivizes voltage stability and penalizes excessive reactive power usage:

$$r = -\frac{1}{|V|} \sum_{i \in V} \left[\alpha_1 |v_i| + \alpha_2 |q_i^{PV}| \right] \quad (9)$$

where:

- $|v_i|$ represents the voltage deviation from the desired level (e.g., 1.0 p.u.).

- $|q_i^{PV}|$ penalizes the reactive power generated by PV inverters.
- α_1, α_2 are weighting factors balancing the trade-off between voltage stability and reactive power usage.

Objective Function: The objective of the reinforcement learning agent is to find a policy π that maximizes the cumulative discounted reward:

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right] \quad (10)$$

where γ is the discount factor, encouraging long-term stability and efficient voltage regulation.

This detailed setup ensures the environment accurately models real-world operational constraints and challenges in a PV-penetrated microgrid.

A.3 EXPERIMENT SETUP AND HYPERPARAMETERS

For the voltage regulation environment, $\alpha_1 = 1$ and $\alpha_2 = 0.1$. For the hyperparameters of both BCQ and CQL, please refer to Table 1. The experiment is implemented on a Windows system with an Intel 13th Gen i7 processor, an NVIDIA GeForce 4060ti GPU, and PyTorch 1.10.

Table 1: Hyperparameters for BCQ and CQL

Hyperparameter	BCQ	CQL
Q-network learning rate	10^{-4}	10^{-4}
Target network update rate (τ)	0.005	0.005
Batch size	64	64
Discount factor (γ)	0.99	0.99
Maximum action perturbation (Φ)	0.05	–
Behavior cloning weight (λ)	0.75	–
VAE latent space dimension	64	–
Number of sampled actions per state	10	10
Temperature parameter (α)	–	0.1
Penalty weight for conservative loss (ζ)	–	0.5
Regularization weight for policy constraint (β)	–	1.0

A.4 DATASET STATISTICS

Table 2 lists several performance metrics for the Expert, Medium, and Poor datasets. These metrics offer insights into the three offline datasets’ variability, quality, and “experience level” (i.e., how much average valuable information each contains).

Table 2: Summary of Reward and Return Statistics for Expert, Medium, and Poor Datasets

Performance Metric	Expert	Medium	Poor
Average Reward	-0.0242	-0.0678	-0.4325
Reward Variance	1.97E-04	6.5995	7.6837
Reward Std. Deviation	0.014	2.5691	2.227
Average Episode Return	-5.7464	-15.8355	-85.3963
Maximum Episode Return	-2.5771	-13.75	-20.349
Return Variance	4.9976	1.44E+3	7.91E+3

A.5 CASE STUDY ON THE IEEE 33-BUS SYSTEM

The IEEE 33-bus system is a widely used test case for distribution networks, representing a radial network structure with distributed loads and generation, which serves as the testbed for evaluating voltage regulation strategies in a PV-penetrated microgrid. Real-world load and PV generation data are used to simulate operational conditions, providing a comprehensive environment for assessing the performance of different offline-RL algorithms.




Figure 4: The IEEE 33-bus system with solar PV penetration.

Network Configuration: - Each load is represented by a blue circle in Fig. 4. - Specific nodes have installed photovoltaic (PV) generators, marked by solar symbols.

Load Data: - The load data is derived from a 3-year dataset of Portuguese electricity consumption, capturing the usage patterns of 232 consumers. - To introduce variability, the load profile is modified with a random perturbation of $\pm 5\%$ on the power factors.

PV Generation Data: - The PV generation data is sourced from Elia, Belgium’s power network operator. - It is adjusted to a 3-minute resolution to align with the control intervals, ensuring consistency with real-world operational conditions.

Simulation Setup: This system provides a rigorous environment for evaluating different voltage regulation strategies under varying load and generation conditions. The realistic load data, PV penetration, and power flow simulations ensure the algorithms are tested in a representative operational setting.

A.6 PERFORMANCE ON OFFLINE DATASETS

This section provides a detailed evaluation of BCQ’s performance on the Expert, Medium, and Poor datasets, with corresponding metrics and visualizations.




Figure 5: Comparison of the VAE total loss across datasets.

Fig. 5 shows the VAE total loss for each dataset: - Expert dataset: Exhibits a predictable and stable pattern, achieving the lowest reconstruction loss due to its high-quality transitions. - Medium dataset: Has moderate reconstruction loss, reflecting its mixed-quality transitions with added noise. - Poor dataset: Displays the highest reconstruction loss, indicating unstructured and variable data, consistent with its low-quality transitions.

Fig. 6 compares voltage regulation performance across the datasets using the following metrics:

- 1) **Totally Controllable Ratio:** The percentage of steps where all buses maintain voltage levels within the desired range. Higher ratios indicate better controllability, with the Expert dataset achieving the highest ratios, followed by Medium and Poor datasets.
- 2) **Average Episode Return:** Measures cumulative rewards per episode. Even with the Poor dataset, BCQ learns effective strategies and achieves reasonable returns.




Figure 6: Voltage regulation performance across three offline datasets.

3) **Average Voltage Deviation:** Represents the deviation of bus voltages from the target level (1.0 p.u.). The Expert dataset yields the lowest deviation, followed by the Medium and Poor datasets.

4) **Average Ratio of Out-of-Control Buses:** The proportion of buses outside the acceptable voltage range. The Expert dataset shows the smallest ratio, while the Poor dataset exhibits the highest ratio, demonstrating the influence of data quality.

These results highlight BCQ’s ability to adapt to varying data quality levels while maintaining acceptable voltage regulation performance. However, its performance is directly influenced by the quality of the offline dataset, with lower-quality datasets leading to reduced controllability and stability.