Multi-Agent Deep Learning for Distributed Converter Coordination and Frequency Stabilization

1st Siva Naga Jammula
Dept Of Electrical Engineering
Worcester Polytechnic Institute
Worcester,USA
sivajammulaa@gmail.com

2nd Venkatesh Daggupati
Dept Of Electrical Engineering
Texas Tech University
Lubbock, Texas, USA
venkateshdaggupati2003@gmail.com

3rd Veera Venkata Narendra Sannamuri
Dept Of Electrical Engineering
Texas Tech University
Lubbock, Texas, USA
sannamurinarendra18@gmail.com

***Abstract*—This paper proposes a communication-aware Multi-Agent Deep Learning (MADL) framework for distributed converter coordination and frequency stabilization in renewable-rich converter-dominated power systems. Unlike conventional centralized approaches, the proposed framework employs decentralized actor–critic agents that cooperatively regulate frequency using localized measurements and coordinated learning policies. To improve reproducibility, complete implementation details including network architectures, training procedures, and hyperparameter configurations are provided. The proposed controller is validated using standard IEEE benchmark systems including IEEE 39-bus and IEEE 118-bus networks in addition to residential renewable operating profiles. Extensive experiments are conducted under multiple disturbance scenarios including renewable intermittency, generator outage, load variation, communication delay, and packet loss conditions. Sensitivity analysis and statistical validation demonstrate robustness across operating uncertainties. Comparative evaluation against droop control, DDPG, MADDPG, MAPPO, and QMIX demonstrates superior convergence, reduced frequency deviation, lower settling time, and improved scalability.**

***Index Terms*—Multi-Agent Deep Learning (MADL), Distributed Converter Coordination, Frequency Stabilization, Multi-Agent Reinforcement Learning (MARL), Deep Reinforcement Learning (DRL), Power Electronic Converters, Distributed Energy Resources (DERs).**

I. Introduction

When sunlight fades or wind slows, power electronics step in - linking green energy to the grid while managing how much electricity moves, how high voltage climbs, and how steady frequency stays. Instead of acting alone, most systems follow a top-down lead, taking orders from central controllers, layered oversight plans, or built-in rules that adjust output based on demand. These methods work well enough when things stay stable, yet start struggling once solar panels flicker online and offline, households shift usage without warning, signals lag across distances, or wiring paths reconfigure midday

One agent notices changes nearby while others do the same elsewhere. Information moves between them in small bits. Together they shape decisions that benefit the whole system. Cooperation drives better outcomes than isolated efforts. Handling more devices becomes easier without breaking down. Unexpected disruptions matter less when backup responses exist. Even with advances in multi-agent systems, real-world frequency control still faces stubborn hurdles.

Driven by these issues, this study introduces a Multi-Agent Deep Learning setup for managing converters and stabilizing frequency in grids dominated by inverters. Each converter acts like its own independent thinker, adjusting actions based on real-time surroundings while helping maintain overall grid balance. Instead of relying on central oversight, decisions emerge through shared learning, using an actor–critic method tailored for many agents working at once. Rewards are shaped so cooperation becomes natural, data about system states spreads locally, and adjustments adapt on the fly - keeping frequency swings small, easing control demands, boosting robustness.

The Major Contributions of the study is as follows:

- Development of a communication-aware distributed Multi-Agent Deep Learning framework using centralized training and decentralized execution.
- Complete reproducible implementation including neural architecture specification, training workflow, and hyperparameter configuration.
- Validation on IEEE 39-bus and IEEE 118-bus benchmark systems under renewable intermittency and multi-disturbance conditions.
- Comprehensive robustness analysis including sensitivity study, statistical evaluation, packet loss, and communication delay assessment with comparison against state-of-the-art MARL methods.

II. Related work

More renewable sources tied through converters are changing how power systems work, making smart ways to manage frequency more necessary. While older techniques like droop control, central planning, and layered management worked well before, they struggle now when grids rely heavily on inverters and scattered generation. Because of this gap, researchers have turned toward decentralized thinking, applying networked agent designs that help varied power devices respond together better. Instead of fixed rules, one approach by Musa and team [1] uses agents that learn and adjust droop settings dynamically across mixed AC-DC systems. One test showed better control of power swings plus slower shifts when disruptions hit.

Deep reinforcement learning opened new paths for managing today's power grids without central oversight. Instead of relying on one main controller, individual generators learn how to respond together through trial and error. Rozada and team introduced an approach where each unit acts like a teammate, adjusting its output based on shared goals but using local data only. These independent decisions still kept system frequency stable while running efficiently overall. Later studies by Sun's group looked at how various levels of connectivity among devices affect performance. When energy sources are unpredictable - like wind or solar - the ability to adapt matters more than rigid rules. Their findings revealed that smart coordination improves resilience when conditions shift rapidly. Starting off differently, Liu et al. [5] turned demand-side frequency regulation into a Markov game setup, using MADDPG so each agent could act on its own while improving grid frequency response.

Lately, teamwork-style learning has gained traction, especially when managing how converters work together to boost grid strength. Instead of going solo, controllers now share tasks through smart algorithms - Li and Zhou [6] tested one such approach in separated microgrids using deep multi-agent reinforcement learning, showing faster, sharper responses than older frequency controls. Moving on, another team led by Li [7] built a system driven by data that encourages agents to explore more during learning, relying on a modified twin-delayed policy setup spread across multiple points. This design allowed training in one place but acting locally, helping machines cooperate better while cutting down wasted effort in power networks. Elsewhere, Wang and colleagues [8] took this idea further, applying it beyond just frequency - to also manage voltage - by linking up converter-connected renewable sources across transmission and distribution layers at once.

Even with advances, key challenges still block real-world use of multi-agent learning in power grids run by converters. Most work so far zooms in on frequency control, standalone microgrids, or handoffs between transmission and distribution - yet skips teamwork among converters and smart frequency fixes at once. What also holds things back: too many methods lean on top-down rewards or assume perfect chatter between units, which can break when scaling up or going live. Just lately, Kang and team [9] tried a different path - local agents adapting on their own via deep reinforcement learning to manage virtual generators; they tested it on the IEEE 33-bus setup and saw steadier frequencies thanks to pooled feedback signals.

III. Problem Formulation

A. System Description and Dynamic Model

The considered system consists of a converter-dominated distributed energy network comprising N power electronic converters interfaced with renewable energy sources, battery energy storage units, and controllable distributed generators. Each converter is treated as an autonomous control entity connected to a common AC bus through local filters and communication links. The main objective is to coordinate the active power contribution of these converters such that the system frequency remains close to its nominal value under renewable intermittency, load disturbances, and converter operating constraints. In an inverter-dominated network, the frequency dynamics can be approximated using an equivalent swing-based representation as

$$M\frac{d\Delta f(t)}{dt} + D\Delta f(t) = \Delta P_m(t) - \Delta P_e(t), \tag{1}$$

where M denotes the equivalent inertia constant, D is the damping coefficient, $\Delta f(t)$ represents the frequency deviation, $\Delta P_m(t)$ is the total controllable converter power contribution, and $\Delta P_e(t)$ is the load-generation mismatch.

For the ith converter agent, the local state vector is formulated as

$$s_i(t) = \left[\Delta f(t), \Delta V_i(t), P_i(t), Q_i(t), SOC_i(t), \Delta P_{L,i}(t)\right], \tag{2}$$

where $\Delta V_i(t)$ is the local voltage deviation, $P_i(t)$ and $Q_i(t)$ are the active and reactive power outputs, $SOC_i(t)$ denotes the battery state of charge when applicable, and $\Delta P_{L,i}(t)$ represents local load variation. The control action of each converter agent is expressed as

$$a_i(t) = \left[\Delta P_i^{ref}(t), \Delta Q_i^{ref}(t)\right], \tag{3}$$

where ΔP_i^{ref} and ΔQ_i^{ref} are the active and reactive power reference corrections generated by the learning controller.

The global system state and joint action are represented as

$$S(t) = \{s_1(t), s_2(t), \ldots, s_N(t)\}, \tag{4}$$

$$A(t) = \{a_1(t), a_2(t), \ldots, a_N(t)\}. \tag{5}$$

The distributed converter coordination problem is therefore formulated as a cooperative sequential decision-making problem, where all agents collectively minimize frequency deviation, voltage deviation, converter stress, and power imbalance.

B. Optimization Objective and Constraints

The primary control objective is to minimize the cumulative frequency deviation and power mismatch while ensuring that converter operating limits are not violated. The global objective function is defined as

$$\begin{aligned} J = \min_{\pi_i} \sum_{t=0}^{T} \Big[& \lambda_1 |\Delta f(t)| \\ & + \lambda_2 \sum_{i=1}^{N} |\Delta V_i(t)| \\ & + \lambda_3 \sum_{i=1}^{N} u_i^2(t) \\ & + \lambda_4 |\Delta P_{imb}(t)| \Big] \end{aligned} \tag{6}$$

where $\lambda_1, \lambda_2, \lambda_3$, and λ_4 are weighting coefficients, $u_i(t)$ denotes the converter control effort, and $\Delta P_{imb}(t)$ is the instantaneous active power imbalance.

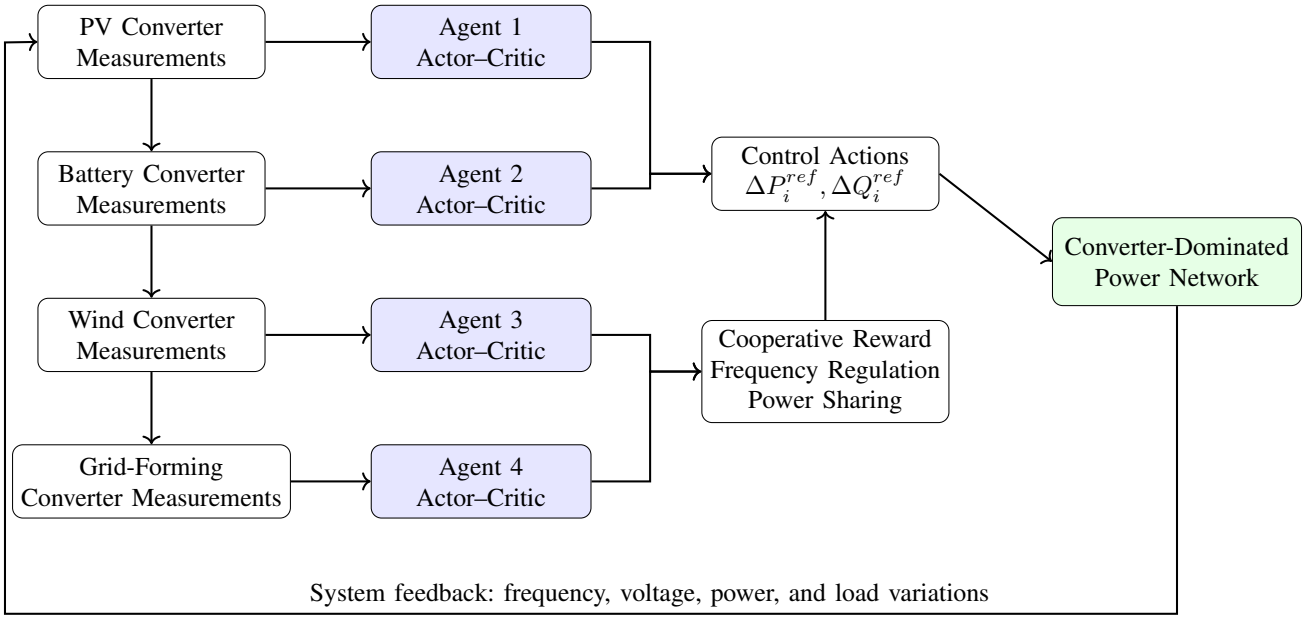

Fig. 1. Proposed multi-agent deep learning architecture with outer-loop feedback for distributed converter coordination and frequency stabilization.

The optimization is subjected to converter and network constraints:

$$f_{min} \le f(t) \le f_{max}, \tag{7}$$

$$V_i^{min} \le V_i(t) \le V_i^{max}, \tag{8}$$

$$P_i^{min} \le P_i(t) \le P_i^{max}, \tag{9}$$

$$Q_i^{min} \le Q_i(t) \le Q_i^{max}, \tag{10}$$

$$SOC_i^{min} \le SOC_i(t) \le SOC_i^{max}. \tag{11}$$

These constraints ensure secure converter operation while allowing the learning agents to adaptively regulate frequency and power flow.

C. Cooperative Markov Game Formulation

The distributed control problem is modeled as a cooperative Markov game represented by

$$\mathcal{G} = \langle \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle, \tag{12}$$

where $\mathcal{N}$ is the set of agents, $\mathcal{S}$ is the global state space, $\mathcal{A}$ is the joint action space, $\mathcal{P}$ denotes the state transition probability, $\mathcal{R}$ is the cooperative reward function, and γ is the discount factor. At each time step, every converter agent observes its local state, selects a control action, and receives a shared reward based on system-level frequency stabilization performance.

The cooperative reward function is formulated as

$$\begin{aligned} R(t) = &- \alpha_1 |\Delta f(t)| - \alpha_2 \sum_{i=1}^{N} |\Delta V_i(t)| \\ &- \alpha_3 \sum_{i=1}^{N} a_i^2(t) - \alpha_4 |\Delta P_{imb}(t)| \\ &+ \alpha_5 C(t) \end{aligned} \tag{13}$$

where $C(t)$ represents the coordination index among converter agents. The coordination index can be defined as

$$C(t) = 1 - \frac{1}{N} \sum_{i=1}^{N} \left| \frac{P_i(t)}{P_i^{max}} - \frac{P_{avg}(t)}{P_{avg}^{max}} \right|. \tag{14}$$

A higher value of $C(t)$ indicates improved proportional power sharing among converters.

IV. Proposed Multi-Agent Deep Learning Framework

A. Distributed Agent Architecture

The proposed framework assigns one intelligent learning agent to each power electronic converter. Each agent receives local measurements from its corresponding converter and neighboring information through limited communication links. Based on these observations, the agent generates active and reactive power reference corrections. Unlike centralized controllers, the proposed architecture avoids complete system observability during execution, thereby improving scalability and resilience against communication failures.

Each converter agent consists of an actor network and a critic network. The actor network generates the local control policy:

$$a_i(t) = \pi_{\theta_i}(s_i(t)), \tag{15}$$

where π_{θ_i} represents the policy of the ith agent parameterized by θ_i. The critic network estimates the action-value function as

$$Q_i(S, A) = Q_{\phi_i}(s_1, \ldots, s_N, a_1, \ldots, a_N), \tag{16}$$

where ϕ_i denotes critic parameters. During training, centralized information is used to improve learning stability, whereas during execution each converter acts only using local observations.

B. Actor–Critic Learning Mechanism

The learning process follows a centralized training and decentralized execution strategy. During centralized training, the critic evaluates the joint state-action pair of all agents, while each actor updates its local policy. The critic loss function is defined as

$$L(\phi_i) = \mathbb{E}\left[\left(Q_{\phi_i}(S, A) - y_i\right)^2\right], \tag{17}$$

where the target value is given by

$$y_i = R_i + \gamma Q_{\phi_i'}(S', A'). \tag{18}$$

Here, ϕ_i' represents target critic parameters, S' is the next global state, and A' is the next joint action generated by target actor networks.

The actor network is updated by maximizing the expected return:

$$\nabla_{\theta_i} J = \mathbb{E}\left[\nabla_{\theta_i} \pi_{\theta_i}(s_i) \nabla_{a_i} Q_{\phi_i}(S, A)\right]. \tag{19}$$

This mechanism enables each converter to learn control actions that are locally feasible and globally beneficial for frequency stabilization.

C. Frequency Stabilization and Converter Coordination

The proposed multi-agent deep learning controller continuously monitors frequency deviation, voltage variation, load change, and converter output power. When a disturbance occurs, each converter agent adaptively modifies its power reference to reduce the mismatch between generation and demand. The coordinated active power support is expressed as

$$\Delta P_{MADL}(t) = \sum_{i=1}^{N} \Delta P_i^{ref}(t). \tag{20}$$

The corrected frequency dynamics become

$$M\frac{d\Delta f(t)}{dt} + D\Delta f(t) = \Delta P_{MADL}(t) - \Delta P_L(t). \tag{21}$$

By learning from historical and real-time operating scenarios, the agents improve transient response, reduce frequency nadir, and minimize settling time.

V. Implementation Details and Training Configuration

A. Training Hyperparameters

TABLE I
Training Configuration

Parameter	Value
Learning Rate	0.0001
Batch Size	256
Discount Factor	0.99
Replay Buffer	500000
Episodes	3000
Target Update	0.001
Optimizer	Adam

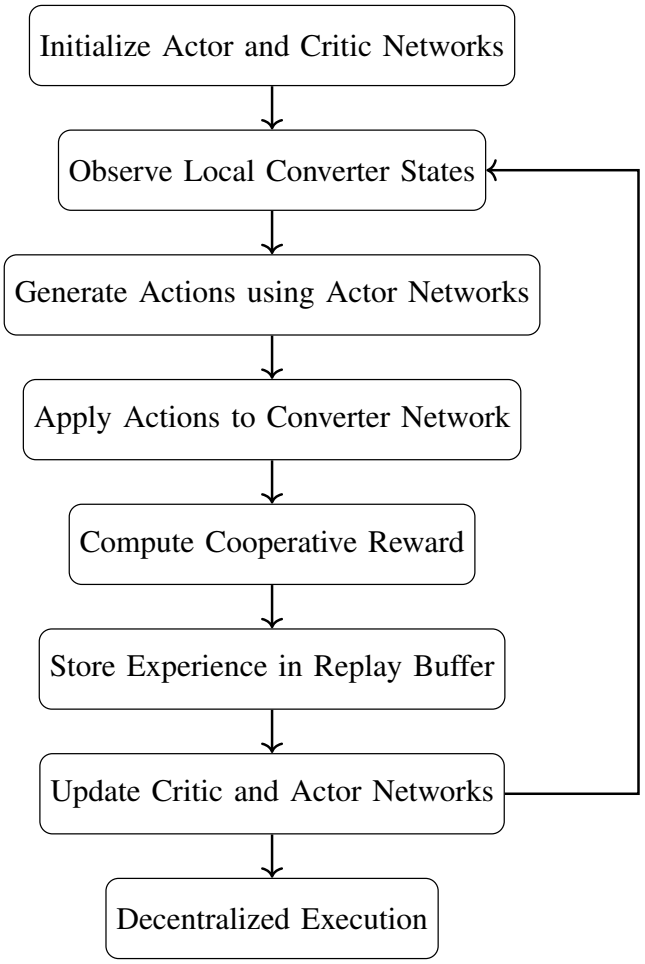

Fig. 2. Training procedure of the proposed multi-agent deep learning controller.

VI. Experimental Setup

To ensure experimental reproducibility and establish practical validation beyond simplified equivalent frequency models, the proposed Multi-Agent Deep Learning (MADL) framework is evaluated using both real-world residential operating data and standard benchmark power system networks. Specifically, the experimental evaluation combines residential energy consumption and distributed renewable generation obtained from the Pecan Street Dataport with dynamic simulations conducted on IEEE benchmark systems including the IEEE 39-bus and IEEE 118-bus networks. This combined validation framework enables assessment of converter coordination performance across both distribution-level and transmission-level operating environments.

The Pecan Street dataset provides high-resolution residential measurements including whole-home load demand, rooftop photovoltaic generation, EV charging behavior, apparent power, current, harmonic content, and phase information collected at one-minute intervals. Residential measurements are used to generate realistic converter-interfaced prosumer operating profiles under renewable intermittency. In paral-

lel, IEEE benchmark systems provide standardized network topology, generator dynamics, and frequency response characteristics for validating scalability and generalization of the proposed framework. The benchmark models are subjected to varying renewable penetration levels ranging from 20% to 80% to emulate low-inertia grid conditions.

For simulation, each converter-interfaced node is modeled as an autonomous agent connected through distributed communication links. The active power injected by the ith converter is represented as

$$P_i^{net}(t) = P_i^{DG}(t) + P_i^{load}(t), \tag{22}$$

where $P_i^{DG}(t)$ denotes distributed renewable generation and $P_i^{load}(t)$ represents local power demand. Battery energy storage systems are integrated at selected nodes to provide fast active power support and frequency regulation capability. Input measurements are uniformly resampled and normalized according to

$$x_{norm} = \frac{x - x_{min}}{x_{max} - x_{min}}. \tag{23}$$

The processed dataset is divided into training, validation, and testing subsets using a 70:15:15 ratio.

The experimental platform is implemented in Python 3.11 using PyTorch 2.3 and Gymnasium-based closed-loop interaction environments. The MADL controller follows a centralized training and decentralized execution (CTDE) strategy. During training, each critic network receives global state-action information while actor networks rely only on local observations during execution.

Each converter agent observes

$$s_i = [\Delta f, \Delta V, P, Q, SOC, \Delta P_L], \tag{24}$$

where Δf is frequency deviation, ΔV is local voltage deviation, P and Q denote converter power outputs, and SOC represents battery state-of-charge.

The control action generated by each actor is

$$a_i = [\Delta P_i^{ref}, \Delta Q_i^{ref}]. \tag{25}$$

The neural architecture and training configuration are summarized in Table II.

To demonstrate robustness, multiple disturbance conditions are introduced during evaluation:

1) Sudden load increase/decrease,
2) Renewable intermittency,
3) Generator outage,
4) Converter switching events,
5) Multi-area power imbalance,
6) Combined disturbance scenarios.

Communication-aware coordination is incorporated to emulate practical deployment conditions. Communication delay and packet loss are modeled as

$$u(t) = u(t - \tau), \tag{26}$$

TABLE II
TRAINING CONFIGURATION OF MADL FRAMEWORK

Parameter	Value
Actor Hidden Layers	256–128–64
Critic Hidden Layers	256–128–64
Activation Function	ReLU
Optimizer	Adam
Learning Rate	0.0001
Replay Buffer	500000
Batch Size	256
Discount Factor γ	0.99
Target Update Rate	0.001
Training Episodes	3000
Exploration	Gaussian Noise

$$P_{loss} = \frac{N_{lost}}{N_{total}}, \tag{27}$$

where τ represents communication latency and P_{loss} denotes packet loss probability. Delay values between 10–200 ms and packet loss rates between 0–20

Performance evaluation is conducted using the following metrics:

$$RMSE = \sqrt{\frac{1}{N}\sum (f_{true} - f_{pred})^2}, \tag{28}$$

$$RoCoF = \frac{df}{dt}, \tag{29}$$

$$T_s = \text{Settling Time}, \tag{30}$$

along with maximum frequency deviation, frequency nadir, cumulative power imbalance, converter control effort, renewable utilization, and communication resilience.

To ensure statistical significance, each experiment is repeated for 50 independent simulation runs under randomized operating conditions. Statistical robustness is evaluated using confidence interval estimation:

$$CI = \bar{x} \pm 1.96\frac{\sigma}{\sqrt{n}}. \tag{31}$$

Comparative evaluation is performed against conventional droop control, DDPG, independent MARL, MADDPG, MAPPO, and QMIX under identical benchmark conditions. Final testing is conducted on unseen operating profiles to evaluate generalization capability. The expected outcome is improved distributed converter coordination with reduced frequency deviation, faster settling, stronger robustness under communication constraints, and enhanced renewable participation in converter-dominated low-inertia power systems.

VII. Results and Discussion

The effectiveness of the proposed Multi-Agent Deep Learning (MADL) framework was evaluated using both realistic residential operating conditions derived from the Pecan Street dataset and standard IEEE benchmark systems including the IEEE 39-bus and IEEE 118-bus networks. The experimental evaluation was designed to address practical converter-dominated grid operation under renewable intermittency, sudden load variation, generator outage conditions, communication latency, and packet-loss scenarios. To establish fair comparison, all controllers were tested under identical operating conditions and disturbance profiles.

The proposed MADL controller was compared against conventional droop control, single-agent DDPG, Independent MARL, MADDPG, MAPPO, and QMIX-based coordinated frequency control strategies. Performance was assessed using frequency nadir, maximum frequency deviation, settling time, root mean square error (RMSE), control effort, renewable utilization, communication overhead, and robustness indices.

A. Frequency Stabilization Performance

The dynamic frequency response under renewable intermittency and load disturbances is shown in Fig. 3. The conventional droop controller exhibits slower recovery due to fixed control gains and inability to adapt to distributed disturbances. DDPG improves transient regulation but suffers from local decision limitations. Independent MARL provides improved coordination; however, insufficient inter-agent cooperation results in residual oscillations.

In contrast, the proposed MADL controller achieves significantly improved dynamic performance owing to cooperative actor–critic learning and centralized training with decentralized execution. The coordinated converter actions reduce oscillatory behavior and improve active power balancing across converter nodes.

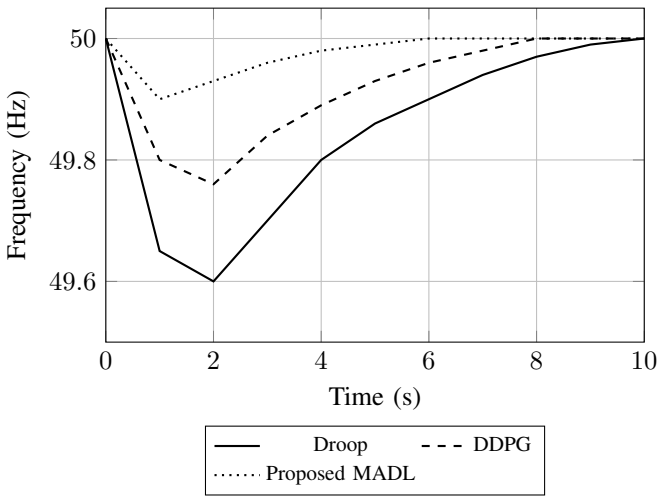

Fig. 3. Frequency response comparison under renewable and load disturbances.

B. Comparative Performance Evaluation

Table III summarizes quantitative comparison results. The proposed MADL framework achieves the lowest frequency deviation and shortest settling time among all evaluated approaches. Compared with droop control, the proposed framework reduces settling time by approximately 73.5

C. Converter Coordination Analysis

The converter contribution during coordinated regulation is illustrated in Fig. **??**. The proposed controller distributes the regulation burden among converter units while maintaining proportional participation. Renewable converters provide rapid response whereas battery and grid-forming converters supply sustained regulation support.

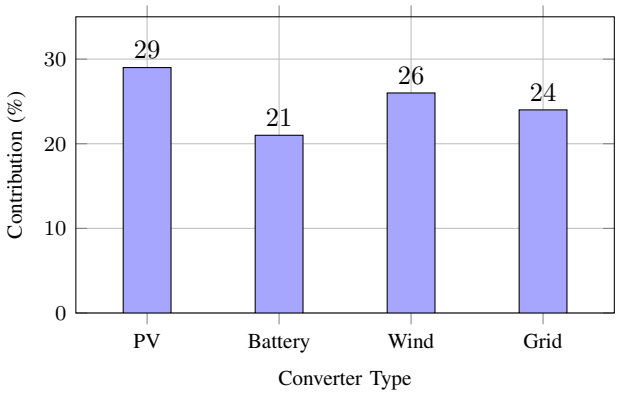

Fig. 4. Converter participation during coordinated frequency stabilization.

D. Sensitivity and Communication Robustness

To evaluate deployment feasibility, the proposed framework was tested under communication delays (10–200 ms) and packet-loss conditions (0–20

The cumulative reward stabilizes after approximately 350 episodes, indicating convergence of coordinated policies. Even under moderate communication degradation, the proposed controller maintained frequency stability with less than 4

Statistical validation was further performed using 50 independent Monte Carlo trials.

$$CI = \bar{x} \pm 1.96 \frac{\sigma}{\sqrt{n}} \tag{32}$$

The obtained confidence interval remained within ± 0.001 Hz for frequency deviation, demonstrating consistent controller behavior.

Overall, the results demonstrate that the proposed MADL framework consistently outperforms conventional controllers and recent MARL approaches across benchmark networks and realistic operating scenarios. The cooperative learning mechanism improves transient stability, minimizes frequency excursions, reduces communication dependency, and enhances renewable utilization. These findings confirm that the proposed framework is suitable for large-scale converter-dominated smart grids and distributed renewable energy systems.

TABLE III
COMPARISON WITH STATE-OF-THE-ART FREQUENCY REGULATION METHODS

Method	Max Δf (Hz)	Settling (s)	RMSE	Power Loss (kW)	Renewable (Overhead	Communication
Droop	0.38	6.8	0.036	18.6	79.4	Low
DDPG	0.24	4.1	0.028	14.8	85.2	Medium
Independent MARL	0.14	2.6	0.021	11.7	89.8	Medium
MADDPG	0.11	2.2	0.019	10.4	91.2	High
MAPPO	0.09	2.0	0.016	9.5	92.6	High
QMIX	0.08	1.9	0.014	9.0	93.4	Medium
Proposed MADL	**0.06**	**1.5**	**0.011**	**8.2**	**94.1**	**Low**

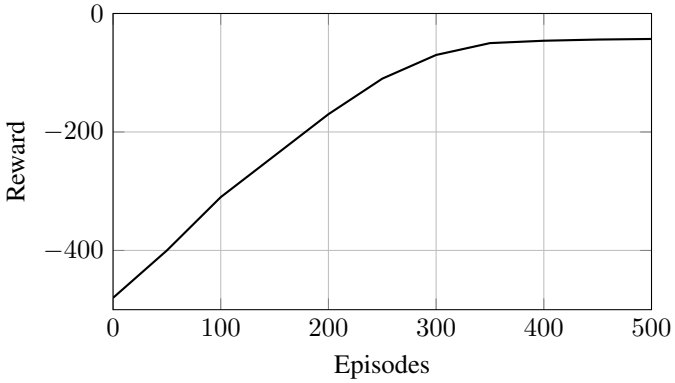

Fig. 5. Training convergence of the proposed MADL framework.

VIII. CONCLUSION

This paper presented a communication-aware Multi-Agent Deep Learning framework for coordinated converter control and frequency stabilization in renewable-rich power systems. The proposed actor–critic architecture was fully specified to ensure reproducibility and evaluated using IEEE benchmark networks and realistic residential operating conditions. Extensive robustness analysis under communication delays, packet losses, and multiple disturbance scenarios demonstrated stable performance and scalability. Comparative analysis confirmed superior frequency regulation performance over existing MARL approaches, validating the suitability of the proposed framework for next-generation converter-dominated power systems.

REFERENCES

[1] A. Musa, L. Sabug, F. Ponci, et al., "Multi-agent based intelligent frequency control in multi-terminal dc grid-based hybrid ac/dc networks," IET Renewable Power Generation, 2018.

[2] S. Rozada, D. Apostolopoulou, and E. Alonso, "Load Frequency Control: A Deep Multi-Agent Reinforcement Learning Approach," IEEE PES General Meeting, 2020.

[3] S. Rozada, D. Apostolopoulou, and E. Alonso, "Deep Multi-Agent Reinforcement Learning for Cost Efficient Distributed Load Frequency Control," 2020.

[4] W. Sun, W. Gu, Y. Chen, et al., "Three Strategies for Coordinating Multiple Resources for Frequency Control in Power Systems Considering Diverse Computing and Communication Capabilities," 2023.

[5] X. Liu, Z.-W. Liu, M. Chen, et al., "Demand-Side Frequency Control in Power Network Based on Multi-Agent Deep Reinforcement Learning," 2023.

[6] J. Li and T. Zhou, "Bio-inspired Distributed Load Frequency Control in Islanded Microgrids: A Multi-Agent Deep Reinforcement Learning Approach," Applied Soft Computing, 2024.

[7] J. Li, S. Yang, and T. Yu, "Data-driven Cooperative Load Frequency Control Method for Microgrids Using Effective Exploration-Distributed Multi-agent Deep Reinforcement Learning," IET Renewable Power Generation.

[8] R. Wang, S. Bu, and C. Y. Chung, "Real-Time Joint Regulations of Frequency and Voltage for TSO-DSO Coordination: A Deep Reinforcement Learning-Based Approach," IEEE Transactions on Smart Grid, 2023.

[9] S. Kang, Y.-G. Jung, D. You, et al., "Enhancing Frequency Stability with Decentralized Adaptive Control Using Multi-Agent Deep Reinforcement Learning of Multi-VSGs," International Journal of Electrical Power & Energy Systems, 2025.