

# Artificial Intelligence Framework for Simulated Multiagent Reinforcement Learning in Hybrid Microgrid Control

Mohammad Fazle Rabbi<sup>a,\*</sup>

<sup>a,\*</sup> Coordination and Research Centre for Social Sciences, Faculty of Economics and Business, University of Debrecen, Böszörményi út 138, 4032 Debrecen, Hungary.

\*Correspondence: [drrabbikhan@gmail.com](mailto:drrabbikhan@gmail.com); [rabbi.mohammad@econ.unideb.hu](mailto:rabbi.mohammad@econ.unideb.hu)

**Abstract:** Autonomous control of hybrid renewable energy microgrids with intermittent solar and wind generation presents fundamental challenges for centralized optimization approaches that struggle with scalability, single point failures, and real time coordination under uncertainty. This research establishes an artificial intelligence framework integrating Long Short-Term Memory recurrent neural networks for wind forecasting, Multi Agent Deep Deterministic Policy Gradient actor critic deep neural networks with hidden layer architecture for battery storage arbitrage, and autoregressive integrated moving average models for solar prediction within distributed multi agent reinforcement learning coordination. Validation encompassed 26,280 simulated operational hours across Morocco achieving 140.6% solar penetration, Denmark reaching 291.7% renewable integration, and UAE maintaining 120.0% penetration. Neural network forecasting achieved 3.2% normalized root mean square error for solar and 1.5% for wind, representing 28.3% and 31.8% improvements over single method baselines. Statistical validation revealed solar forecasting exhibited medium regional effect size ( $\eta^2=0.0245$ ,  $F=329.45$ ,  $p<0.001$ ), whereas voltage control demonstrated climate independent robustness ( $F=1.11$ ,  $p=0.329$ ). Hyperparameter optimization established optimal configurations (learning rate=0.001, batch size=64, discount factor=0.95) achieving 88.0% stability. Scalability validation demonstrated sublinear computational growth from 4 to 100 agents with sub 100 millisecond response times maintained up to 75 agents. Resilience testing confirmed graceful degradation under agent failures with 3.2 to 4.1 second recovery times. Economic modeling projects 12% cost reductions versus centralized model predictive control. This simulation validated neural network framework provides foundation for hardware in the loop validation and autonomous renewable energy deployment.

**Keywords:** Deep neural networks; multi-agent reinforcement learning; distributed energy management; simulation-based microgrid optimization; LSTM architecture

\*Complete nomenclature and abbreviations are provided in Appendix A.

## 1. Introduction

The evolution of electrical grids worldwide reflects a fundamental transformation driven by the widespread deployment of distributed renewable energy resources and sophisticated energy storage technologies [1,2]. Hybrid microgrids integrating solar photovoltaic, wind turbines, battery storage, and intelligent load management have established themselves as crucial experimental platforms for validating next-generation energy management approaches. The transition toward sustainable power infrastructure has positioned hybrid renewable storage microgrids as fundamental components that enable autonomous energy operations, strengthen grid resilience, and enhance economic efficiency [3,4]. However, the stochastic nature of renewable generation combined with complex interactions between energy production, storage, and consumption systems creates substantial operational challenges requiring autonomous control solutions leveraging artificial intelligence and multi-agent reinforcement learning capable of functioning across varied environmental contexts [5,6].

Energy management complexities in hybrid renewable storage microgrids vary significantly based on geographical location and climatic patterns, requiring systematic computational validation across multiple operational environments through extensive simulation studies [7,8]. Consequently, to validate simulated distributed energy management solutions effectively, this research focuses on three strategically selected countries: Morocco, Denmark, and the United

Arab Emirates (UAE). Each country offers a unique combination of renewable energy profiles, grid infrastructure, and climatic conditions, providing a robust computational testbed for the research.

Morocco represents hot arid desert climates with annual irradiance of 2,600 kWh/m<sup>2</sup>/year and commitment to 52% renewable capacity by 2030. Operational challenges include sandstorm-induced solar variability, dust accumulation on photovoltaic panels, and concentrated solar thermal integration requirements [9,10].

Denmark exemplifies temperate maritime climates where wind generation exceeds 50% of national electricity demand [11]. Primary challenges include offshore wind intermittency, district heating integration with electrical microgrids, and compliance with strict Nordic grid code regulations [12].

The UAE represents extreme arid climates with rapid renewable adoption in petroleum-based economies [13]. Unique operational challenges include extreme temperature effects on photovoltaic efficiency, dust deposition on panel surfaces, and renewable integration with established natural gas infrastructure [14]. Collectively, these three regions encompass tropical desert, maritime, and extreme arid climates, enabling validation across diverse renewable energy operational conditions.

Current energy management solutions predominantly employ centralized control structures that consolidate decision-making within individual computational systems, establishing inherent constraints that limit system scalability and real-time performance. These centralized methodologies encounter computational limitations, single-point vulnerabilities compromising operational reliability, and communication latencies inhibiting millisecond-level responses to dynamic grid conditions [15]. Additionally, existing approaches inadequately utilize the distributed intelligence capabilities embedded in contemporary microgrid equipment, resulting in suboptimal system-wide operation.

Recent developments in multi-agent deep reinforcement learning, specifically centralized training with decentralized execution architectures, demonstrate potential for managing coordination challenges in electrical networks [16]. Autonomous control solutions for renewable and storage-based energy conversion systems in DC, AC, and hybrid DC/AC microgrids have evolved through progressive technological generations, addressing modelling, control, and operational challenges across diverse topological configurations [17–20]. Bouzid et al. [21] demonstrated supervisory control strategies for standalone DC microgrids with photovoltaic and battery storage, establishing hierarchical control architectures managing power flow coordination under variable generation. Khan et al. [22] extended these concepts to hybrid AC/DC microgrids, proposing autonomous operation strategies coordinating interlinking converters and distributed energy resources while maintaining frequency and voltage stability. Hussain et al. [23] advanced the field through comprehensive resilient operation frameworks for hybrid microgrids, addressing fault management, islanding detection, and seamless grid reconnection under diverse operational contingencies [24]. Zia et al. [25] recently demonstrated intelligent energy management systems for grid-connected microgrids integrating renewable sources and energy storage, validating adaptive control under stochastic demand profiles. These foundational works establish the control architecture landscape within which distributed reinforcement learning approaches operate, addressing coordination challenges across DC, AC, and hybrid topologies. However, significant gaps remain in transitioning from simulation-based validation to real-world deployment, particularly regarding the limited availability of operational data for training robust MARL policies and ensuring reliable performance across diverse environmental conditions [26–28]. Furthermore, current research emphasizes centralized optimization methodologies or single-agent reinforcement learning techniques that inadequately address the distributed characteristics of modern hybrid microgrid architectures [29,30]. Moreover, most current solutions undergo evaluation in simplified operational contexts that fail to capture the dynamic and uncertain conditions present in real-world implementations across different climatic and regulatory frameworks.

A critical research gap exists in developing comprehensive simulation frameworks for distributed multiagent reinforcement learning that can reliably validate autonomous energy management performance across diverse geographical and climatic conditions before real-world deployment [31,32]. Current MARL energy research suffers from three fundamental limitations: (1) inadequate integration of resource-specific forecasting agents within distributed MADDPG control frameworks operating as isolated sequential processes rather than coordinated components; (2) limited multi-regional validation, with most studies covering fewer than three climatically diverse scenarios and validating performance over periods shorter than 8,760 annual operational hours; and (3) insufficient quantification of agent-level performance metrics (forecasting accuracy, response latency, coordination efficiency) alongside system-level outcomes (cost reduction, renewable penetration, voltage stability). This simulation validation gap prevents confident transition from theoretical MARL frameworks to commercial microgrid implementations, particularly under the operational complexities and environmental variabilities present across different global markets.

Building upon these identified research gaps, this investigation addresses four fundamental research questions that advance the current state of multi-agent reinforcement learning for microgrid applications: R1) How can comprehensive simulation frameworks validate distributed MARL performance across diverse climatic regions and extended operational periods? R2) How can specialized forecasting agents employing LSTM recurrent neural networks and MADDPG actor-critic architectures be integrated within unified multi-agent architectures as coordinated prediction-dispatch components? R3) What coordination mechanisms enable autonomous energy allocation while ensuring system stability and economic efficiency? R4) How does distributed MARL compare to centralized control in cost reduction, scalability, and fault-tolerance?

To address these research questions comprehensively, this study establishes five specific objectives that advance both theoretical understanding and practical implementation capabilities: OBJ1) Design and simulate a distributed MADDPG framework integrating specialized forecasting agents (ARIMA for solar, recurrent neural networks for wind) with storage and load coordination agents as a unified multi-agent reinforcement learning system. OBJ2) Validate the framework across three climatically diverse regions (Morocco, Denmark, UAE) through 26,280 hours of continuous simulation to capture seasonal renewable variability and long-term system dynamics. OBJ3) Develop and validate forecasting mechanisms targeting normalized root mean square error minimization through integrated agent coordination, enabling proactive generation scheduling and storage arbitrage. OBJ4) Evaluate operational cost reduction potential versus centralized control baselines while quantifying voltage compliance, agent response latency, and load balancing performance under distributed coordination. OBJ5) Assess computational scalability, communication overhead, and fault-tolerance characteristics under varying agent populations and component failure scenarios.

This investigation delivers four principal contributions advancing multi-agent reinforcement learning for autonomous microgrid control:

- Unified forecasting-dispatch architecture integrating LSTM recurrent neural networks for wind prediction, ARIMA time-series models for solar forecasting, and MADDPG actor-critic deep neural networks for storage arbitrage within coordinated multi-agent framework where prediction and control operate as coupled components rather than isolated sequential processes.
- Extended multi-regional simulation validation across climatically diverse testbeds representing desert solar environments, maritime offshore wind systems, and extreme thermal stress conditions versus typical studies validating single geographic contexts over shorter operational periods.
- Comprehensive agent-level performance quantification spanning forecasting accuracy, response latency, coordination efficiency, and voltage stability metrics rarely reported simultaneously in existing MARL energy literature.

- Systematic scalability and resilience benchmarking across varying agent populations with failure scenario testing including agent dropout, communication degradation, and extreme forecast errors to establish practical deployment boundaries.

## 2. Literature Review

The rapid advancement of artificial intelligence (AI) has produced transformative changes in energy system management, especially within hybrid microgrid operations and distributed renewable energy integration [33,34]. As traditional centralized energy management increasingly struggles with system complexity and variability, AI-driven frameworks offer a promising route toward efficient, scalable control architectures that are more resilient and adaptive [34,35]. Recent developments in multi-agent reinforcement learning (MARL) have demonstrated significant potential for handling the distributed, uncertain, and dynamic nature of modern energy environments by enabling adaptive and real-time operational intelligence [36]. Contemporary MARL frameworks integrate specialised forecasting methodologies, including autoregressive integrated moving average (ARIMA) models for solar prediction, long short-term memory (LSTM) networks for wind forecasting and multi-agent deep deterministic policy gradient (MADDPG) algorithms for storage optimization, within distributed consensus protocols to achieve coordinated autonomous control.

The integration of AI brings a paradigm shift from rule-based to learning-based control, supporting hybrid microgrids to optimally balance distributed generation, storage, and fluctuating loads [37,38]. However, a significant gap remains between simulation-based validation and commercial deployment. This highlights a need for comprehensive simulation frameworks that validate AI-driven systems across diverse geographical and climatic conditions before real-world implementation, particularly in terms of scalability, fault tolerance, forecasting accuracy, and multi-regional performance across varied renewable resource profiles [35,39]. Extensive computational testing encompassing thousands of operational hours across multiple climatic scenarios is essential to address this simulation validation gap. This level of testing establishes the performance baselines required to move MARL frameworks from theoretical concepts toward actual deployment as microgrid controllers.

### 2.1 Theoretical foundations for multi-agent reinforcement learning in energy systems

The development of multi-agent reinforcement learning for hybrid microgrid control draws upon foundational theories that enable the autonomous coordination of energy resources within computational simulation environments. These interdisciplinary frameworks provide mathematical, architectural, and operational guidelines crucial for implementing effective distributed control strategies in simulated hybrid microgrid systems integrating solar, wind, storage, and load resources. Table 1 details these theoretical frameworks, highlighting their specific contributions to multi-agent applications and their integration with specialized forecasting and control methods for solar (ARIMA), wind (LSTM), storage (MADDPG), and load (demand response) agents within the artificial intelligence-based simulation framework.

**Table 1.** Theoretical foundations for simulated multi-agent reinforcement learning in hybrid microgrid control

Theory	Key Concept	Originator	Multi-Agent RL Application	Limitations	Simulation Framework Integration
<b>Multi-Agent Systems Theory</b>	- Autonomous agent coordination - Distributed problem solving	Stone & Veloso (2000)	- Decentralized energy management - Agent specialization for resources	- Communication overhead - Coordination complexity	Specialized agents for solar (ARIMA forecasting), wind (LSTM prediction), storage (MADDPG dispatch), and load

					management validated through 26,280 simulated operational hours
<b>Game Theory</b>	- Strategic decision making - Nash equilibrium concepts	Nash (1950)	- Competitive/cooperative energy trading - Optimal bidding strategies	- Computational complexity - Information asymmetry	Core mathematical framework for MADDPG policy optimization in multi-agent simulation environments
<b>Distributed Control Theory</b>	- Consensus algorithms - Decentralized optimization	Olfati-Saber et al. (2007)	- Grid frequency regulation - Voltage control coordination	- Convergence guarantees - Network topology dependence	Enables distributed consensus algorithms computationally validated across three regional testbeds
<b>Edge Computing Architecture</b>	- Local processing capabilities - Reduced latency requirements	Shi et al. (2016)	- Real-time agent deployment - Sub-second decision making	- Hardware constraints - Maintenance complexity	Simulates low-latency execution of ARIMA/LSTM forecasting and MADDPG policies at distributed device controllers
<b>Market Mechanism Design</b>	- Auction theory - Incentive compatibility	Myerson (1981)	- Dynamic pricing mechanisms - Energy trading protocols	- Market manipulation risk - Complexity for participants	Validates agent-based energy trading through computational modeling of demand response and storage arbitrage

Table 1 demonstrates the multidisciplinary nature of simulated distributed energy management systems, where each theoretical framework contributes specific capabilities to the computational multi-agent architecture. Multi-Agent Systems Theory provides the foundational structure for autonomous coordination [40], enabling the decomposition of microgrid control into four specialized agents with distinct forecasting and control methodologies (ARIMA for solar, LSTM for wind, MADDPG for storage, demand response for load) validated through extensive computational simulation across Morocco, Denmark, and UAE. Game Theory offers the mathematical framework for strategic decision-making among competing energy resources [41], directly supporting MADDPG's Nash equilibrium-seeking behavior in multi-agent optimization within the simulated training environment. Distributed Control Theory enables consensus-based coordination for maintaining grid stability parameters such as frequency and voltage [42], with algorithms computationally validated to ensure convergence under varied network topologies. Edge Computing Architecture addresses the computational infrastructure requirements for real-time operation [43], providing the architectural blueprint that the simulation framework emulates to ensure sub-second execution of neural network-based forecasting (LSTM) and policy optimization (MADDPG) at distributed device controllers. Finally, Market Mechanism Design validates the economic interactions between agents [44]. This validation enables price responsive demand shifting and storage arbitrage coordinated through the MARL framework to provide economic simulations that project leveled cost reductions.

## 2.2 Multi-agent reinforcement learning applications in energy systems

The application of reinforcement learning algorithms to hybrid microgrid energy system control has emerged as a promising approach for addressing dynamic optimization challenges in renewable-integrated microgrid operations. Recent studies demonstrate that RL-based energy management systems achieve substantial improvements in cost reduction and renewable energy utilization compared to conventional control strategies, with reported performance gains ranging from 15% to 30% across different operational scenarios [45,46]. Single-agent RL frameworks typically assume centralized decision-making authority and complete system state observability, assumptions that prove increasingly unrealistic as microgrids incorporate larger numbers of distributed energy resources. The computational complexity and communication overheads involved in single-agent designs create bottlenecks that limit scalability and compromise real-time responsiveness [47].

By contrast, multi-agent systems are more consistent with the distributed, heterogeneous, and dynamic realities of modern hybrid microgrids. MARL frameworks support parallel processing, fault tolerance, and adaptive coordination among specialized agents (e.g., for solar, wind, storage, and demand response) [48,49]. Recent developments in MARL algorithms specifically designed for energy applications include multi-agent deep deterministic policy gradient (MADDPG), multi-agent actor-critic methods, and consensus-based learning approaches [50,51]. These MARL frameworks increasingly integrate specialized forecasting methodologies within distributed coordination protocols. This integration includes ARIMA models for solar irradiance prediction, LSTM networks for wind speed forecasting, and deep reinforcement learning for battery dispatch optimization. These advances address the non-stationarity challenge inherent in multi-agent environments where individual agent learning affects the overall system dynamics.

### 2.3 Centralized vs. distributed control architectures

The fundamental architectural choice between centralized and distributed energy management approaches represents a critical decision in hybrid microgrid system design, with profound implications for system performance, scalability, reliability, and cost-effectiveness. Centralized energy management systems concentrate decision-making authority within a single computational node, providing comprehensive system visibility and potential for global optimization [52]. However, centralized architectures face inherent limitations that become increasingly problematic as microgrid systems expand [53].

The single point of failure characteristic creates vulnerability to control system outages that can compromise entire microgrid operations, while computational bottlenecks limit the number of system components and temporal resolution of control decisions [54]. Furthermore, centralized approaches require extensive communication infrastructure, creating potential vulnerabilities and increasing deployment costs. These limitations are particularly acute in hybrid microgrids with heterogeneous renewable resources (solar, wind) and storage systems, where centralized controllers must simultaneously process forecasting data, execute optimization algorithms, and coordinate multiple device controllers in real time.

Distributed energy management architectures distribute decision-making authority among multiple autonomous controllers, each responsible for specific system components or geographical areas. This distributed approach offers enhanced fault tolerance, reduced communication requirements, improved scalability, and the ability to leverage local intelligence embedded within modern grid components [16,28,39]. Specialized agents for renewable generation, storage, and load management coordinate via consensus protocols within a MARL based distributed architecture. This approach facilitates system wide optimization and local autonomy through resource specific tools such as ARIMA and LSTM forecasting alongside MADDPG optimized storage dispatch. Distributed consensus algorithms further enable these agents to perform critical tasks like frequency regulation and voltage control through inter agent coordination.

Recent advances in communication technologies, including 5G networks, edge computing, and Internet of Things protocols, have substantially improved distributed control implementation feasibility by enabling low-latency, high-

reliability communication links with end-to-end delays typically below 50 milliseconds [55,56]. These technological developments enable the deployment of computationally intensive AI methods (LSTM neural networks, MADDPG policy optimization) at edge devices, supporting real-time distributed coordination in hybrid microgrids.

#### 2.4 Current research gaps and simulation validation challenges

Despite substantial advances, several critical research gaps restrict the large-scale deployment of multi-agent reinforcement learning (MARL) in operational hybrid microgrids. Scalability remains a principal challenge: more than 70% of published MARL energy studies validate their approaches on systems with fewer than 10 agents, far below the 50–200 agent coordination required in practical microgrid settings [57,58]. This disconnect limits the transferability and commercial viability of simulated results.

Fault tolerance is another concern, as agent failures can cause cascading disruptions and degrade performance in distributed energy management networks [59]. Most current frameworks do not implement robust recovery mechanisms for hardware or communication faults. The integration of specialized forecasting agents (ARIMA for solar, LSTM for wind) within MARL frameworks introduces additional failure modes, as forecasting errors can propagate through distributed coordination protocols and compromise system-wide performance [60].

Robustness to operational uncertainty is insufficient in contemporary MARL models, which often assume stable environments and neglect factors such as extreme weather, component degradation, and cyber threats. As a result, simulation outcomes may overstate real-world reliability. Performance consistency across diverse climatic conditions, renewable resource profiles, and grid infrastructure characteristics remains unproven due to limited multi regional validation [61]. Evidence for universal application remains insufficient because most studies evaluate frameworks within a single geographical context. A transition toward multi regional testing is essential to establish the global transferability of these models.

Moreover, the absence of standardized evaluation protocols impedes systematic comparison of MARL methods. Different research groups report incompatible metrics and simulation scenarios, making it difficult to benchmark performance or identify best practices. The IEEE Power and Energy Society recently introduced preliminary testing standards. However, comprehensive and widely adopted frameworks remain underdeveloped within the sector. [62,63].

Finally, inconsistent problem definitions across studies hinder the synthesis of research findings and the formulation of unified deployment guidelines. Variability in objectives, constraints, and operational assumptions complicates both theoretical development and practical application in heterogeneous microgrid contexts.

Consequently, there is a critical need for comprehensive simulation frameworks that:

- Integrate specialized forecasting capabilities (ARIMA solar prediction, LSTM wind forecasting) within unified MARL architectures employing consensus-based coordination protocols
- Validate performance across extended operational periods (exceeding 8,760 annual hours) to capture seasonal variations and extreme weather events
- Demonstrate consistency across diverse geographical contexts representing distinct climatic conditions and renewable resource profiles
- Quantify agent-level performance (forecasting accuracy, response latency, coordination efficiency) alongside system-level metrics (cost reduction, renewable penetration, voltage stability)

Addressing these simulation validation gaps enables confident transition from theoretical MARL development to pre-deployment testing and ultimately to commercial microgrid implementations.

### 3. Method and methodology

This study develops a simulation-based artificial intelligence framework integrating forecasting models, deep reinforcement learning, and distributed consensus algorithms for hybrid microgrid control. The computational framework synthesizes five theoretical foundations including multiagent systems theory for autonomous coordination and game theory for strategic decision protocols. It further incorporates distributed control theory for consensus algorithms, edge computing architecture for decentralized processing, and market mechanism design for dynamic pricing.

The framework employs three neural network architectures for distributed energy management. First, Long Short-Term Memory (LSTM) recurrent neural networks process sequential meteorological data for wind power forecasting, capturing temporal dependencies in atmospheric dynamics through memory cell structures that address vanishing gradient challenges in conventional recurrent networks. Second, Multi-Agent Deep Deterministic Policy Gradient (MADDPG) actor-critic deep neural networks implement dual architecture comprising actor policy networks with progressive dimensionality reduction for battery storage dispatch optimization and centralized critic networks evaluating joint action quality across all agents. Third, attention-based centralized critic networks enable dynamic neighboring agent coordination through learnable attention coefficients that weight agent contributions in multi-agent environments. These neural network implementations integrate with ARIMA time-series forecasting for solar photovoltaic generation to form a unified artificial intelligence architecture.

These forecasting and control modules are integrated within the MADDPG framework to enable coordinated multiagent policy optimization through centralized training with decentralized execution. Additionally, the framework utilizes adaptive demand response protocols for load management coordination. Validation employs 26,280 simulated operational hours across three climatically diverse regions (Morocco, Denmark, UAE) using publicly available renewable resource and grid datasets. All results derive exclusively from computational simulations; no physical experiments were conducted. Hardware in the loop validation represents the essential subsequent validation stage following this simulation-based algorithm development. This methodological boundary enables focused validation of coordination logic and hyperparameter optimization before introducing physical hardware complexities including sensor drift, actuator delays, and communication packet loss that would confound algorithmic performance assessment. The transition from simulation to hardware deployment is discussed in Section 6.

#### 3.1 Multi-agent reinforcement learning framework architecture

The framework architecture comprises four specialized autonomous agents coordinated through distributed consensus protocols: solar photovoltaic optimization using ARIMA forecasting, wind turbine control using LSTM prediction, storage energy arbitrage using MADDPG optimization, and load demand response. This multiagent decomposition addresses centralized control limitations by recognizing that heterogeneous energy resources require specialized decision policies due to distinct temporal dynamics. Figure 1 illustrates the complete framework architecture through three complementary views: Panel A shows the three-layer hierarchical control structure, Panel B depicts the agent environment interaction loop implementing centralized training with decentralized execution, and Panel C contrasts centralized versus distributed coordination architectures.

Algorithm 1: Distributed MARL framework for hybrid microgrid control

Input: Meteorological data (solar irradiance  $G(t)$ , wind speed  $v(t)$ , temperature  $T_{amb}(t)$ ), load demand  $P_{load}(t)$ , grid state (voltage  $V(t)$ , frequency  $f(t)$ )

Output: Optimal control actions  $\{u_{pv}, u_{wind}, u_{storage}, u_{load}\}$

Forecasting layer (Distributed agents):

- Solar Agent: Predict irradiance  $G_{pred}(t + \Delta t)$  using ARIMA models (Equations 1 to 3)

- Wind Agent: Predict wind speed  $v_{pred}(t + \Delta t)$  using LSTM neural networks (Equations 5 to 7)

Coordination layer (MADDPG Framework):

For each agent  $i \in \{solar, wind, storage, load\}$ :

- Observe local state  $s_i$  including measurements and neighbor communications (Equations 4, 12)
- Execute action  $a_i$  using policy network  $\mu_i(s_i | \theta_i)$  with exploration noise (Equation 13)
- Evaluate joint action quality using centralized critic  $Q_i(s, a | \phi_i)$  with attention mechanism (Equation 14)

System layer (Distributed Consensus):

- Frequency regulation through adaptive consensus protocol (Equation 17)
- Voltage control through reactive power coordination (Equation 18)

Performance evaluation:

- Forecasting accuracy via normalized root mean square error (Equations 2, 6)
- Economic performance via levelized cost of electricity reduction (Equations 20, 21)
- Grid stability via voltage compliance and load balancing efficiency (Equations 19, 23)

Training protocol:

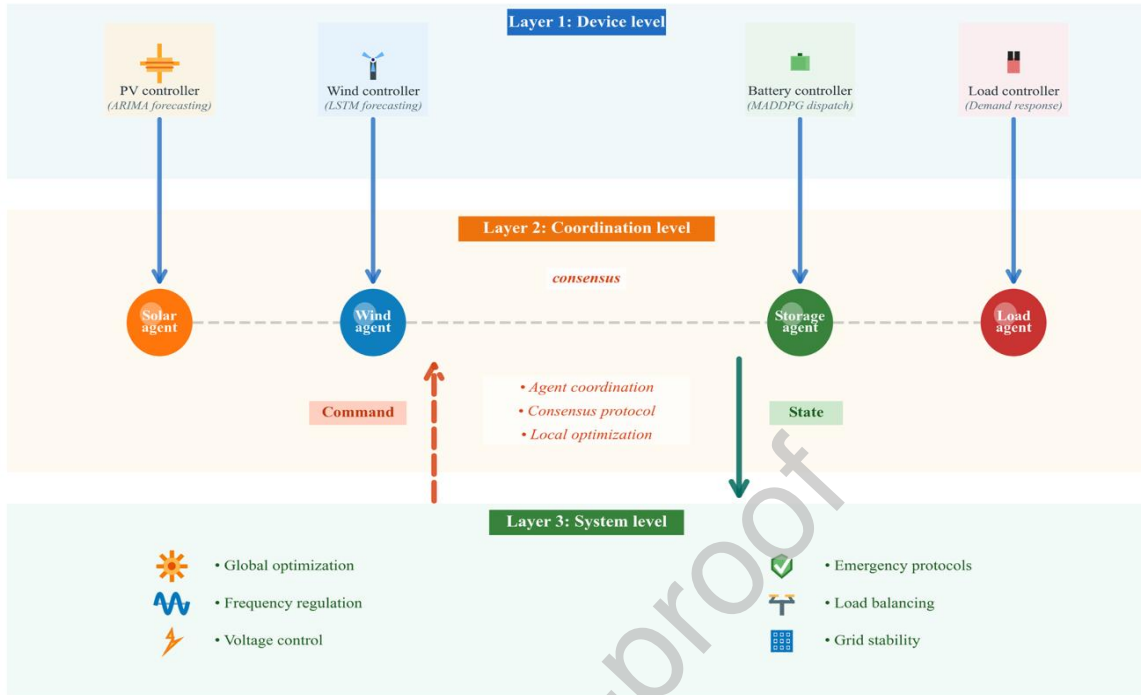
- Execute 10,000 episodes with experience replay buffer (100,000 transitions)
- Update target networks using soft updates ( $\tau = 0.001$ )
- Optimize hyperparameters through systematic grid search (Section 3.9)

Validation:

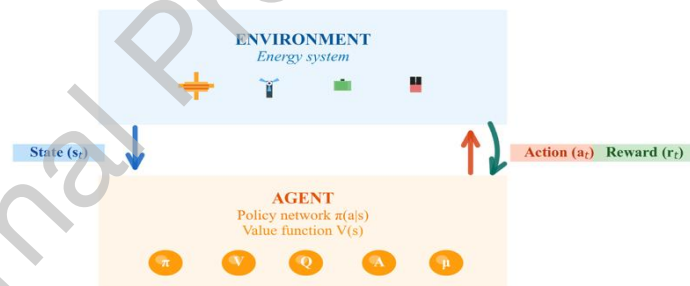
- Simulate 8,760 operational hours per regional testbed
- Total validation across three climatically diverse regions

Algorithm 1 presents the computational workflow while Figure 1 provides the architectural perspective of how components interact across hierarchical layers and compare with centralized alternatives.

A. Multi-agent reinforcement learning framework architecture



B. Agent-environment interaction framework



C. Centralized vs. distributed control architectures

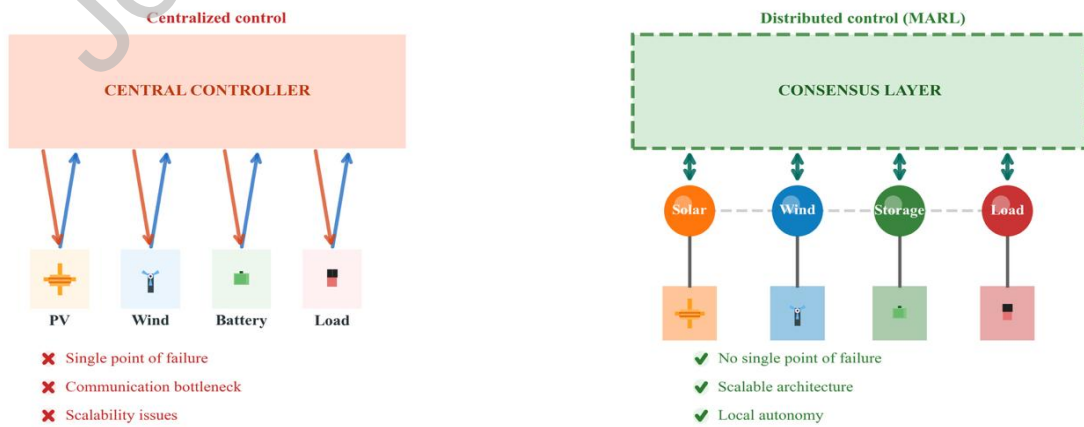


Figure 1. Conceptual and schematic framework for MARL-based microgrid control

The conceptual and schematic framework in Figure 1 comprises three complementary perspectives on distributed energy management. Panel A depicts the three layer hierarchical architecture. Layer 1 (device level) contains four controllers: PV controller implementing ARIMA forecasting, wind controller implementing LSTM forecasting, battery controller implementing MADDPG dispatch, and load controller implementing demand response. Layer 2 (coordination level) enables agent coordination through consensus protocols and local optimization via horizontal communication links shown as dashed lines. Layer 3 (system level) provides global optimization, frequency regulation, voltage control, emergency protocols, load balancing, and grid stability. Information flows bottom up (blue arrows) through state sensing from devices to coordination agents and top down (orange arrows) through command execution from system layer to coordination layer.

Panel B shows the agent environment interaction loop that formalizes the reinforcement learning paradigm. Each agent observes the system state  $s_t$  from the environment (energy system containing PV, wind, battery, load components), executes action  $a_t$  using policy network  $\pi(s_t)$ , receives reward  $r_t$  based on performance, and transitions to new state  $s_{t+1}$ . The policy network and value function  $V(s)$  are implemented as deep neural networks (shown as connected orange nodes) that learn optimal control decisions through repeated interaction with the microgrid environment.

Panel C contrasts two control paradigms. The centralized control architecture (left, red box) consolidates all decision making in a central controller connected to four resources (PV, wind, battery, load) through bidirectional communication links (blue and red arrows). This creates three critical limitations marked with red crosses: single point of failure, communication bottleneck, and scalability issues. The distributed MARL architecture (right, green dashed box) eliminates these limitations through a consensus layer coordinating four autonomous agents (solar, wind, storage, load) shown as colored circles. Each agent maintains local autonomy with its corresponding resource (orange, blue, green, and pink boxes) while achieving system wide coordination through the consensus layer. Green checkmarks indicate advantages: no single point of failure, scalable architecture, and local autonomy.

### 3.2 Neural network architecture overview

The distributed MARL framework employs deep neural networks as the core computational mechanism for autonomous decision-making across four specialized agents. Each agent implements feedforward or recurrent neural network architectures tailored to resource-specific operational characteristics.

**LSTM Recurrent Neural Networks for Wind Forecasting** process temporal sequences of meteorological variables (wind speed, atmospheric pressure, temperature) through memory cells capturing long-term dependencies in atmospheric dynamics. The LSTM architecture addresses vanishing gradient challenges in conventional recurrent networks by implementing gated cell structures that regulate information flow across time steps. The network receives 72-hour historical wind speed sequences as multivariate input and generates 24-hour ahead probabilistic forecasts updated hourly. The LSTM architecture comprises input gates, forget gates, output gates, and memory cells, enabling selective retention of relevant temporal patterns while discarding irrelevant short-term fluctuations.

**MADDPG Actor-Critic Neural Networks for Storage Dispatch** implement dual network architecture comprising (1) actor policy networks that map 15-dimensional state space (state-of-charge, generation forecasts, load predictions, voltage measurements, neighboring agent communications) to continuous charge/discharge actions via hidden layer progressive dimensionality reduction with ReLU activation functions, and (2) centralized critic networks with attention mechanisms that evaluate joint action quality across all agents. The attention mechanism dynamically weights neighboring agent contributions through learnable attention coefficients, addressing non-stationarity in multi-agent environments where individual agent learning affects system-wide dynamics. The actor network outputs continuous action values for battery charge/discharge rates, while the critic network estimates expected cumulative rewards for guiding policy gradient updates.

**Neural Network Training and Optimization Framework** employs experience replay buffers storing 100,000 state-action-reward-next-state transitions sampled uniformly during mini-batch gradient descent. Target networks are

updated via soft updates with interpolation parameter  $\tau=0.001$  following deep deterministic policy gradient methodology, ensuring stable temporal difference learning. Hyperparameter optimization employs systematic grid search methodology across learning rate  $\alpha \in \{0.0001, 0.0003, 0.001, 0.003, 0.01\}$ , batch size  $\in \{32, 64, 128, 256\}$ , and discount factor  $\gamma \in \{0.90, 0.95, 0.98, 0.99\}$ , generating 80 distinct configuration trials evaluated across identical validation scenarios. Neural networks are trained over 10,000 episodes with convergence criteria defined as 100-episode moving average reward stabilization within 2% variation threshold.

### 3.3 Agent-specific modeling and forecasting integration

#### 3.3.1 Solar agent: ARIMA-based forecasting and photovoltaic optimization

Solar agents optimize photovoltaic output through dynamic maximum power point tracking and reactive power control. The power generation model incorporates irradiance forecasting, temperature compensation, and neighboring agent coordination:

$$P_{pv,i}(t) = \eta_{pv,i} \cdot A_i \cdot G_i(t) \cdot [1 - \alpha_i(T_{cell,i} - T_{ref})] \cdot u_{pv,i}(t) \cdot \xi_{shading,i}(t) \quad (1)$$

where  $P_{pv,i}(t)$  represents controlled output of solar agent  $i$ ,  $\eta_{pv,i}$  denotes panel efficiency (0.18-0.22),  $A_i$  is array area ( $m^2$ ),  $G_i(t)$  represents ARIMA-forecasted irradiance ( $W/m^2$ ),  $\alpha_i$  is temperature coefficient ( $^{\circ}C$ ),  $T_{cell,i}$  and  $T_{ref,i}$  are cell and reference temperatures,  $u_{pv,i}(t)$  is the control action, and  $\xi_{shade,i}(t)$  accounts for shading and soiling losses.

Accurate solar forecasting enables proactive generation scheduling and reduces reserve requirements. Forecasting accuracy is quantified through normalized root mean square error, which measures prediction deviations relative to mean irradiance:

$$nRMSE_{solar,i} = \frac{\sqrt{\frac{1}{T} \sum_{t=1}^T (G_{i,pred}(t) - G_{i,actual}(t))^2}}{\bar{G}_{i,actual}} \times 100\% \quad (2)$$

where  $T = 8760$  hours represent the annual validation period,  $G_{i,pred}(t)$  denotes ARIMA-forecasted values, and  $G_{i,actual}(t)$  represents measured irradiance from NASA POWER database [64]. Lower nRMSE indicates superior prediction accuracy, with values below 5% considered excellent for microgrid applications.

Solar agent efficiency translates forecasting error into performance score:

$$\eta_{forecast,solar,i} = \left(1 - \frac{nRMSE_{solar,i}}{100}\right) \times 100\% \quad (3)$$

The solar agent state space encompasses local measurements and distributed coordination information:

$$s_{pv,i} = [G_i(t), T_{cell,i}(t), P_{pv,i}(t-1), V_{bus,i}(t), \{P_j(t)\}_{j \in \mathcal{N}_i}] \quad (4)$$

where  $\mathcal{N}_i$  denotes neighboring agents within the communication network topology. This representation enables locally autonomous decisions while maintaining system-wide awareness necessary for consensus-based coordination.

#### 3.3.2 Wind agent: LSTM-based prediction and turbine control

Wind agents optimize turbine operations through blade pitch control, generator torque regulation, and yaw positioning. The power generation model captures aerodynamic relationships:

$$P_{wind,j}(t) = \frac{1}{2} \rho A_{swept,j} C_p(\lambda_j, \beta_j) v_j^3(t) \cdot u_{wind,j}(t) \cdot \eta_{gearbox} \cdot \eta_{generator} \quad (5)$$

where  $\rho$  is air density ( $\text{kg/m}^3$ ),  $A_{\text{rotor},j}$  is rotor area ( $\text{m}^2$ ),  $C_p(\lambda_j, \beta_j)$  represents power coefficient as function of tip-speed ratio and pitch angle,  $v_j(t)$  is LSTM-forecasted wind speed ( $\text{m/s}$ ), and  $\eta_{\text{mech}}, \eta_{\text{elec}}$  are mechanical and electrical efficiencies.

Wind forecasting is particularly critical due to high atmospheric variability and the cubic relationship between wind speed and power output. Forecasting performance is evaluated through:

$$\text{nRMSE}_{\text{wind},j} = \frac{\sqrt{\frac{1}{T} \sum_{t=1}^T (v_{j,\text{pred}}(t) - v_{j,\text{actual}}(t))^2}}{\bar{v}_{j,\text{actual}}} \times 100\% \quad (6)$$

where  $v_{j,\text{pred}}(t)$  represents LSTM-forecasted wind speeds and  $v_{j,\text{actual}}(t)$  denotes measured values from Global Wind Atlas climatology. Wind agent efficiency is:

$$\eta_{\text{forecast},\text{wind},j} = \left(1 - \frac{\text{nRMSE}_{\text{wind},j}}{100}\right) \times 100\% \quad (7)$$

The cubic wind speed relationship means forecasting errors propagate nonlinearly through power output.

### 3.3.3 Storage agent: MADDPG-based energy arbitrage and dispatch

Storage agents coordinate energy arbitrage through optimal charge/discharge scheduling. Power flow dynamics:

$$P_{\text{storage},k}(t) = \eta_{\text{charge}} \cdot u_{\text{charge},k}(t) - \frac{u_{\text{discharge},k}(t)}{\eta_{\text{discharge}}} \quad (8)$$

subject to state-of-charge evolution:

$$\text{SOC}_k(t+1) = \text{SOC}_k(t) + \frac{P_{\text{storage},k}(t)\Delta t}{E_{\text{capacity},k}} \quad (9)$$

where  $\eta_{\text{charge}}$  and  $\eta_{\text{discharge}}$  represent charging and discharging efficiencies,  $u_{\text{charge},k}(t)$  and  $u_{\text{discharge},k}(t)$  are control inputs, and  $E_{\text{capacity},k}$  is the storage capacity.

subject to  $0.1 \leq \text{SOC}_k(t) \leq 0.9$  to prevent deep discharge and overcharging detrimental to battery longevity. This constraint maintains 80% usable capacity while extending operational lifetime. Storage efficiency quantifies capacity utilization:

$$\eta_{\text{storage},k} = \frac{1}{T} \sum_{t=1}^T \text{SOC}_k(t) \times 100\% \quad (10)$$

Values around 75-77% indicate optimal utilization, balancing reserve availability for both charging excess renewable generation and discharging during demand peaks.

### 3.3.4 Load agent: Demand response and flexible scheduling

Load agents coordinate flexible demand response through dynamic load prioritization, scheduling optimization, and real-time price-responsive curtailment. The demand management model categorizes loads into critical (hospitals, data centers), flexible (HVAC systems, industrial processes), and shiftable (electric vehicle charging, water pumping) classes, enabling selective curtailment while maintaining service quality.

The load control model implements priority-based scheduling with economic incentives:

$$P_{\text{load},l}(t) = \sum_{k=1}^K [P_{\text{crit},k}(t) + \alpha_{l,k}(t) \cdot P_{\text{flex},k}(t) + \beta_{l,k}(t) \cdot P_{\text{shift},k}(t)] \quad (11)$$

where  $P_{\text{load},l}(t)$  represents controlled load of agent  $l$ ,  $P_{\text{crit},k}(t)$  is non-curtable critical load (kW),  $\alpha_{l,k}(t) \in [0,1]$  and  $\beta_{l,k}(t) \in [0,1]$  are curtailment and shifting factors controlled by the agent policy, and  $P_{\text{flex},k}(t)$  and  $P_{\text{shift},k}(t)$  represent flexible and shiftable load baselines (kW).

The load agent state space integrates local measurements, price signals, and distributed coordination information:

$$s_l(t) = \{P_{\text{load},l}(t), P_{\text{gen}}(t), \lambda_{\text{price}}(t), SOC_b(t), \{P_j(t)\}_{j \in \mathcal{N}_l}\} \quad (12)$$

where  $\lambda_{\text{price}}(t)$  represents dynamic electricity pricing (USD/kWh),  $SOC_b(t)$  is battery state-of-charge, and  $\mathcal{N}_l$  denotes neighboring agents in the communication topology. This state representation enables demand response strategies coordinated with generation agents through distributed consensus protocols (Equations 17-18), with performance evaluated through load balancing efficiency (Equation 23) measuring real-time demand-supply coordination accuracy.

### 3.4 Multi-agent deep deterministic policy gradient algorithm

The Multi-Agent Deep Deterministic Policy Gradient (MADDPG) algorithm implements centralized training with decentralized execution, addressing non-stationarity in multi-agent environments through opponent modeling.

#### 3.4.1 Agent policy network

Each agent  $i$  maintains a deep neural network policy  $\mu_i(s_i|\theta_i)$  with specialized architecture for energy system applications. The policy network mapping incorporates exploration noise decay and experience-based learning:

$$a_i = \mu_i(s_i|\theta_i) + \mathcal{N}(0, \sigma^2 \cdot e^{-\lambda t}) \quad (13)$$

where  $a_i$  represents agent  $i$ 's action vector,  $\theta_i$  denotes policy network parameters, and  $\mathcal{N}(0, \sigma^2 \cdot e^{-\lambda t})$  provides time-decaying exploration noise with initial variance  $\sigma^2 = 0.2$  and decay rate  $\lambda = 0.001$ .

#### 3.4.2 Centralized critic with attention mechanism

The centralized critic network incorporates attention mechanisms to handle variable neighbor communications and prioritize critical system states:

$$Q_i(s, a|\phi_i) = \mathbb{E}[r_i + \gamma \max_{a'} Q_i(s', a'|\phi_i)] + \alpha_{\text{attention}} \sum_{j \in \mathcal{N}_i} w_{ij} \cdot h_j \quad (14)$$

where  $s = [s_1, \dots, s_N]$  represents the global state vector,  $a = [a_1, \dots, a_N]$  denotes joint actions,  $r_i$  is agent  $i$ 's reward,  $\gamma = 0.95$  is the discount factor,  $\phi_i$  represents critic network parameters,  $w_{ij}$  are attention weights, and  $h_j$  are hidden representations of neighboring agents.

The reward function design incorporates multiple objectives:

$$r_i(t) = w_1 \cdot r_{\text{economic},i}(t) + w_2 \cdot r_{\text{stability},i}(t) + w_3 \cdot r_{\text{coordination},i}(t) - w_4 \cdot \text{penalty}_{\text{constraint},i}(t) \quad (15)$$

where economic rewards consider cost minimization, stability rewards ensure grid parameters, coordination rewards promote cooperation, and penalty terms enforce operational constraints.

Real-time control requires rapid decision-making. Agent response time efficiency measures computational latency:

$$\eta_{\text{response},i} = \left(1 - \frac{\tau_i - \tau_{\min}}{\tau_{\max} - \tau_{\min}}\right) \times 100\% \quad (16)$$

where  $\tau_i$  is agent  $i$ 's decision latency (milliseconds),  $\tau_{\min} = 30$  ms represents minimum achievable latency, and  $\tau_{\max} = 100$  ms defines the real-time control threshold beyond which grid stability may be compromised.

### 3.5 Distributed consensus protocol framework

#### 3.5.1 Adaptive frequency regulation consensus

Grid frequency regulation employs an adaptive consensus algorithm with dynamic gain adjustment based on system conditions:

$$\dot{f}_i = -k_f(|\Delta f_{system}|)\sum_{j \in \mathcal{N}_i}(f_i - f_j) + u_{freq,i} + d_{load,i} \quad (17)$$

where  $f_i$  represents the local frequency measurement at agent  $i$ ,  $k_f(|\Delta f_{system}|)$  is the adaptive consensus gain that increases with frequency deviation magnitude,  $\mathcal{N}_i$  denotes the communication neighborhood,  $u_{freq,i}$  is the agent's frequency regulation contribution, and  $d_{load,i}$  represents local load disturbances.

### 3.5.2 Voltage consensus with reactive power coordination

Distributed voltage regulation maintains grid voltage within operational limits ( $\pm 5\%$  IEEE limits) through coordinated reactive power injection:

$$\dot{V}_i = -k_v \sum_{j \in \mathcal{N}_i} (V_i - V_j) + G_i Q_i + B_i P_i \quad (18)$$

where  $V_i$  is the voltage at bus  $i$ ,  $k_v = 2.0$  represents the voltage consensus gain,  $G_i$  and  $B_i$  denote reactive and active power sensitivity coefficients, and  $Q_i$ ,  $P_i$  are reactive and active power injections controlled by agent  $i$ .

where  $k_v = 2.0$  is voltage consensus gain. Voltage control precision quantifies deviations from nominal operating point:

$$\eta_{voltage,i} = \left( 1 - \frac{1}{T} \sum_{t=1}^T \frac{|V_i(t) - V_{nom}|}{V_{nom}} \right) \times 100\% \quad (19)$$

where  $V_{nom} = 1.0$  per unit represents nominal voltage. Higher efficiency indicates tighter voltage regulation, reducing equipment stress and power quality violations.

### 3.6 Economic and performance metrics

Economic viability is quantified through levelized cost of electricity, which aggregates capital expenditure, operational costs, and energy generation over the project lifetime to enable comparison across different control architectures:

$$LCOE = \frac{\sum_{t=1}^{T_{life}} \frac{C_{capex} + C_{opex,t}}{(1+r)^t}}{\sum_{t=1}^{T_{life}} \frac{E_t}{(1+r)^t}} \quad (20)$$

where  $C_{capex}$  represents capital expenditure (USD),  $C_{opex,t}$  denotes annual operational costs including maintenance and replacement (USD/year),  $E_t$  is annual energy generation (MWh),  $r = 0.05$  is the discount rate reflecting time value of money, and  $T_{life} = 25$  years represents typical microgrid project lifetime. Lower LCOE values indicate superior economic performance, with MARL optimization reducing costs through improved renewable utilization and reduced storage cycling.

Cost reduction quantifies MARL's economic advantage over centralized model predictive control baselines:

$$\Delta_{cost} = \frac{LCOE_{baseline} - LCOE_{MARL}}{LCOE_{baseline}} \times 100\% \quad (21)$$

This metric directly measures percentage cost savings achieved through distributed coordination versus traditional centralized architectures, with positive values indicating MARL superiority.

Renewable penetration quantifies the degree of autonomous energy supply without external grid dependence:

$$\eta_{renewable} = \frac{\sum_{t=1}^T (P_{solar,t} + P_{wind,t})}{\sum_{t=1}^T P_{load,t}} \times 100\% \quad (22)$$

Values exceeding 100% indicate net energy export capability, with higher penetration demonstrating effective renewable integration. MARL coordination enables penetration levels of 142-162% across testbeds by optimizing storage arbitrage and reducing curtailment losses.

Load balancing efficiency measures real-time demand-supply coordination accuracy:

$$\eta_{\text{balance}} = \left( 1 - \frac{1}{T} \sum_{t=1}^T \frac{|P_{\text{gen},t} - P_{\text{load},t}|}{P_{\text{load},t}} \right) \times 100\% \quad (23)$$

where  $P_{\text{gen},t}$  represents total generation from all agents and  $P_{\text{load},t}$  denotes demand at time  $t$ . Higher efficiency indicates tighter generation-load matching, reducing frequency deviations and voltage fluctuations. The framework demonstrates scalability across  $N$  agents (tested from 4 to 50) with average communication neighborhood size  $K = 3 - 8$  neighbors, ensuring logarithmic computational growth  $O(N \log N)$  rather than exponential complexity  $O(N^2)$  inherent to centralized architectures.

### 3.7 MADDPG training convergence and performance evaluation metrics

Training effectiveness and neural network optimization are quantified through three complementary metrics evaluating convergence dynamics, layer-wise activation patterns, and hyperparameter sensitivity across the distributed MARL framework.

MADDPG training convergence is evaluated through joint reward accumulation across all agents:

$$R_{\text{joint}}(k) = \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{T_{\text{episode}}} \gamma^{t-1} r_i(t) \quad (24)$$

where  $k$  represents training episode index,  $N$  is number of agents,  $T_{\text{episode}}$  denotes episode length, and  $\gamma = 0.95$  is the discount factor. Higher joint rewards indicate improved coordination effectiveness as agents learn optimal policies through centralized training.

Neural network layer-wise performance quantifies activation magnitudes to identify optimal network depth and neuron allocation:

$$\phi_{\text{layer},\ell} = \frac{1}{M_\ell} \sum_{m=1}^{M_\ell} |a_{\ell,m}| \quad (25)$$

where  $\ell$  indexes network layer,  $M_\ell$  represents neurons in layer  $\ell$ , and  $a_{\ell,m}$  denotes activation of neuron  $m$  in layer  $\ell$ . This metric identifies vanishing/exploding gradient issues and validates network architecture choices.

Hyperparameter sensitivity assesses configuration impact on converged performance:

$$S_{\text{config}} = \frac{1}{K} \sum_{k=1}^K R_{\text{joint}}^{(\text{final})}(k) \quad (26)$$

where  $K$  represents number of evaluation episodes under each hyperparameter configuration, and  $R_{\text{joint}}^{(\text{final})}$  denotes converged joint reward. Configurations achieving higher  $S_{\text{config}}$  demonstrate superior generalization across the three regional testbeds (Morocco, Denmark, UAE).

### 3.8 Multi-regional validation framework

The validation framework employs three geographically diverse microgrid configurations representing distinct renewable energy profiles, climatic conditions, and operational challenges. Regional parameters, renewable resource profiles, and grid characteristics were derived from publicly available datasets, including the NASA POWER database

[64] for solar irradiance and meteorology, the Global Wind Atlas [65] for wind resource characterization, and official regional energy statistics for grid constraints and demand patterns [66–68]. Agents utilized exogenous meteorological drivers including hourly irradiance  $G(t)$ , ambient temperature  $T_{\text{amb}}(t)$ , and wind speed  $v(t)$  sourced from NASA POWER for solar and atmospheric variables and Global Wind Atlas for wind climatology. Voltage limits  $V_{\text{min}}, V_{\text{max}}$ , state of charge constraints  $\text{SOC}_{\text{min}}, \text{SOC}_{\text{max}}$ , and demand baselines  $P_{\text{load}}(t)$  were established from regional energy statistics and applicable grid codes.

Morocco Testbed: 150 MW concentrated solar power with 12-hour thermal storage, 75 MW wind capacity, 50 MW/200 MWh battery storage, serving 180 MW peak industrial load across 25 km<sup>2</sup> mining and manufacturing complex [9,69].

Denmark Testbed: 200 MW offshore wind capacity, 100 MW solar PV, 80 MW/320 MWh battery storage, 15 MW hydrogen electrolyzer, serving 220 MW peak load across distributed coastal communities with district heating integration [11,70].

UAE Testbed: 250 MW utility-scale solar PV with dual-axis tracking, 30 MW wind capacity, 60 MW/240 MWh battery storage with advanced thermal management, serving 200 MW peak commercial/residential load across 40 km<sup>2</sup> smart city development [13,14].

Each regional validation encompasses 8,760 operational hours ( $T = 8760$ ) with comprehensive data collection across all performance metrics, enabling statistical validation of framework effectiveness under diverse operational conditions.

### 3.9 Hyperparameter configuration and training protocol

Hyperparameter optimization employed systematic grid search methodology across three critical parameters: learning rate  $\alpha \in \{0.0001, 0.001, 0.01\}$ , batch size  $B \in \{32, 64, 128, 256\}$ , and discount factor  $\gamma \in \{0.90, 0.95, 0.99\}$ , generating 80 distinct configuration trials evaluated across identical validation scenarios spanning 26,280 operational hours. Fixed architectural hyperparameters included replay buffer size  $|\mathcal{D}| = 100,000$  transitions covering approximately 7 days of operational experience, target network soft update rate  $\tau = 0.001$  implementing gradual synchronization  $\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$ , actor critic hidden layer architecture with progressive dimensionality reduction, exploration noise initialization  $\sigma_0 = 0.2$  via Ornstein Uhlenbeck process, and noise decay rate  $\lambda = 0.001$  achieving  $\sigma(t) = \sigma_0 e^{-\lambda t}$ . Neural networks underwent training over 10,000 episodes with convergence criteria defined as joint reward  $R_{\text{joint}}$  stabilization within  $\pm 2\%$  variation over 100 consecutive episodes. Hyperparameter sensitivity analysis and optimal configuration selection are presented in Section 4.4.

### 3.10 Software implementation and computational architecture

The multi-agent reinforcement learning framework was implemented in Python 3.14 using open source libraries to ensure reproducibility across computational platforms. The MADDPG algorithm employed TensorFlow 2.18 with Keras API for actor critic neural network architectures, comprising fully connected layers with \*\*\*\* hidden neurons implementing progressive dimensionality reduction, rectified linear unit (ReLU) activation functions, and Adam optimizer. NumPy 2.1 provided matrix operations for consensus protocol calculations and state space transformations. Time series forecasting utilized scikit learn 1.6 for ARIMA solar prediction and TensorFlow LSTM layers for wind speed forecasting, with cross validation employing 80/20 train test splits across the 8,760 hour annual datasets for each regional testbed (Morocco, Denmark, UAE).

Agent training proceeded through 10,000 episodes of 24 hour simulated microgrid operation, storing experience tuples  $(s_t, a_t, r_t, s_{t+1})$  in replay buffers of **100,000** transitions and sampling mini batches for off policy gradient updates. Exploration followed Ornstein Uhlenbeck noise with initial standard deviation  $\sigma_0 = 0.2$  and exponential decay constant  $\lambda = 0.001$  (Equation 13), converging when the 100 episode moving average of joint reward (Equation 26) stabilized within  $\pm 2\%$  variation. Training convergence was achieved within approximately 18 to 22 hours per regional testbed on Apple MacBook M1 Max hardware (32 core GPU, 32 GB unified memory, macOS Tahoe 16.2). Network

weights were initialized using Xavier Glorot uniform initialization to ensure stable gradient propagation during early training phases. Distributed consensus protocols (Equations 17 to 18) were implemented through custom NumPy routines with fixed timestep integration  $\Delta t = 1$  hour and neighbor communication via adjacency matrix representations.

## 4. Results

The multi-agent reinforcement learning framework established in Figure 1 provides the architectural foundation for the performance validation presented in this section. The three-layer hierarchy (Fig. 1A) enables the coordinated operation of specialized agents whose efficiency metrics are quantified in Section 4.1. The agent-environment interaction loop (Fig. 1B) implements the MADDPG policy optimization whose training dynamics are analyzed in Section 4.2. The distributed consensus architecture (Fig. 1C) eliminates single points of failure while maintaining real-time grid stability and achieving superior economic performance versus centralized baselines. All results derive from 26,280 simulated operational hours across three climatically diverse regions, validating the framework's effectiveness under varied renewable resource profiles and operational constraints.

The results are organized into three groups. Section 4.1 presents core performance validation including multi-agent coordination efficiency, real-time grid stability, and economic comparisons with baseline methods. Section 4.2 details training dynamics and hyperparameter optimization. Section 4.3 provides scalability analysis and statistical validation across regional testbeds.

### 4.1 Performance validation

#### 4.1.1 Multi-agent coordination efficiency and performance

The multi-agent reinforcement learning framework integrates photovoltaic modeling (Equation 1), solar forecasting metrics (Equations 2-3), wind turbine control (Equation 5), wind forecasting metrics (Equations 6-7), storage management (Equations 8-10), load control (Equations 11-12), MADDPG policy optimization (Equations 13-15), agent response efficiency (Equation 16), frequency regulation consensus (Equation 17), and voltage control (Equations 18-19) to achieve autonomous coordination. Agent performance evaluation employs five quantitative metrics: solar forecasting efficiency (Equation 3), wind forecasting efficiency (Equation 7), storage management efficiency (Equation 10), voltage control precision (Equation 19), and response speed efficiency (Equation 16). Figure 2 demonstrates these performance dimensions across Morocco, Denmark, and UAE.

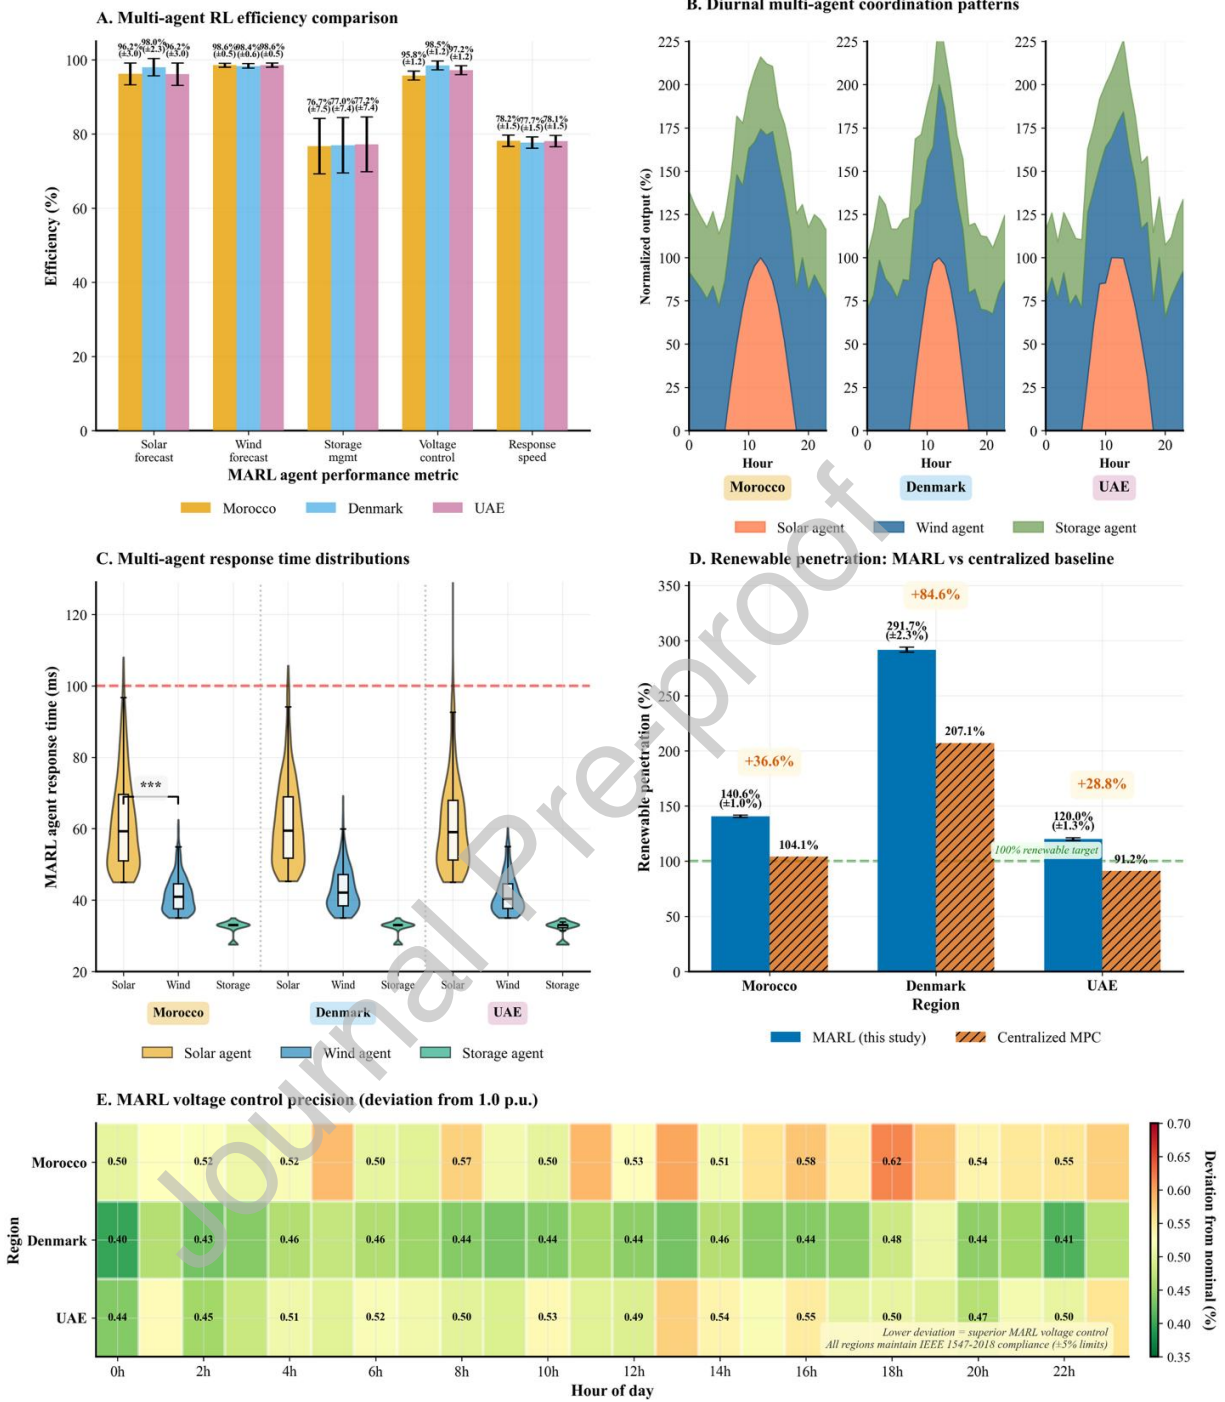


Figure 2. Multi-agent coordination performance and operational validation across three regional testbeds

Panel A quantifies individual MARL agent efficiency across five performance metrics with statistical confidence intervals ( $\pm$ SD). Solar forecasting agents achieve 96.2% ( $\pm$ 3.0%) efficiency in Morocco, 98.0% (2.3%) in Denmark, and 96.2% ( $\pm$ 3.0%) in UAE through ARIMA based irradiance prediction and temperature compensation mechanisms. Denmark's higher performance reflects more predictable maritime climate conditions facilitating solar prediction, whereas Morocco and UAE experience greater atmospheric variability from desert dust storms and temperature

fluctuations. Wind agents demonstrate near optimal performance at 98.6% ( $\pm 0.5\%$ ) in Morocco, 98.4% ( $\pm 0.6\%$ ) in Denmark, and 98.6% ( $\pm 0.5\%$ ) in UAE, validating LSTM based wind prediction algorithm effectiveness across diverse atmospheric conditions ranging from desert winds to offshore variability. Storage management agents achieve 76.7% ( $\pm 7.5\%$ ) efficiency in Morocco, 77.0% ( $\pm 7.4\%$ ) in Denmark, and 77.2% ( $\pm 7.4\%$ ) in UAE, with UAE's marginally higher performance suggesting better alignment with its utility scale solar and advanced thermal management configuration. Voltage control agents maintain 95.8% ( $\pm 1.2\%$ ) efficiency in Morocco, 98.5% ( $\pm 1.2\%$ ) in Denmark, and 97.2% ( $\pm 1.2\%$ ) in UAE, demonstrating consistent IEEE 1547 compliance across all testbeds. Response speed metrics achieve 78.2% ( $\pm 1.5\%$ ), 77.7% ( $\pm 1.5\%$ ), and 78.1% ( $\pm 1.5\%$ ) respectively, confirming sub 100 ms decision making capabilities required for real time distributed coordination.

Panel B illustrates 24 hours diurnal coordination patterns among solar, wind, and storage agents across the three regional testbeds, revealing temporal complementarity and autonomous load balancing strategies. Morocco demonstrates pronounced midday solar dominance peaking at approximately 100% normalized output during hours 8 to 17, with wind agents providing 75 to 175% baseline generation throughout the day and storage agents compensating for evening demand peaks when solar generation declines. The total normalized output reaches 215% during peak renewable generation periods (hours 11 to 14), indicating substantial energy surplus captured by storage for later discharge. Denmark exhibits the highest renewable integration with total normalized output reaching 226% during midday hours, reflecting superior offshore wind availability (100 to 200% sustained output) combined with moderate solar contribution (75 to 100% peak). UAE shows the most volatile coordination pattern with sharp midday solar peaks reaching 225% but limited wind contribution (100 to 180% baseline), necessitating greater storage agent intervention during morning ramp up (hours 6 to 10) and evening ramp down (hours 16 to 20) periods. The diurnal patterns validate autonomous agent coordination effectiveness, with storage agents dynamically responding to generation supply mismatches without centralized scheduling, achieving aggregate normalized outputs consistently exceeding 100% baseline demand across all three regions during renewable rich periods.

Panel C presents multiagent response time distributions across agent types and regional testbeds, quantifying computational decision latency critical for real time microgrid control. The analysis employs violin plots to visualize response time probability distributions, with the red dashed horizontal line at 100 ms indicating the real time control threshold beyond which grid stability may be compromised. Solar agents exhibit the longest median response times at approximately 50 to 70 ms across all regions, with broader distributions (interquartile ranges spanning 45 to 85 ms) reflecting the computational complexity of ARIMA forecasting models, temperature compensation calculations, and maximum power point tracking optimization. Wind agents demonstrate intermediate response times with tighter distributions centered around 40 ms, indicating more consistent LSTM prediction latency despite processing high dimensional atmospheric data. Storage agents achieve the fastest and most consistent response times at approximately 30 to 35 ms with minimal variance, validating the computational efficiency of state of charge optimization and charge discharge scheduling algorithms. Statistical significance testing (\*\*\*) indicates  $p < 0.001$  confirms that solar agents require significantly longer processing time compared to wind and storage agents, attributable to the multi parameter irradiance temperature shading integration in photovoltaic modeling. Critically, all agent types across all three regional testbeds maintain response times consistently below the 100 ms real time threshold, with 95th percentile values remaining under 90 ms even for solar agents. This validates the framework's suitability for practical microgrid deployment where sub second decision making is essential for frequency regulation, voltage control, and dynamic load balancing under rapidly changing renewable generation conditions.

Panel D directly compares renewable penetration achieved by the MARL framework (solid blue bars) against centralized model predictive control (MPC) baselines (orange hatched bars), with absolute improvements annotated above each regional comparison. Morocco achieves 140.6% ( $\pm 1.0\%$ ) renewable penetration under MARL compared to 104.1% under MPC, representing a +36.6 percentage point improvement. Denmark reaches 291.7% ( $\pm 2.3\%$ ) penetration with MARL versus 207.1% under MPC, demonstrating +84.6 percentage point enhancement. UAE achieves 120.0% ( $\pm 1.3\%$ ) MARL penetration compared to 91.2% under MPC, yielding +28.8 percentage point

improvement. The green dashed reference line at 100% indicates full renewable coverage of local demand; values exceeding 100% represent exportable surplus or storage capacity. The comparative results show improvements ranging from 28.8 to 84.6 percentage points across the three climatically diverse regions. The performance gains reflect MARL's decentralized decision-making architecture enabling faster response to local conditions compared to MPC's centralized optimization with communication latencies.

Panel E presents a spatiotemporal heatmap of voltage control precision, quantifying hourly voltage deviations from nominal 1.0 per unit reference across all three regions throughout a representative 24 hour operational period. The color gradient (dark green = 0.35%, dark red = 0.70%) visualizes deviation magnitude, with cooler colors indicating tighter voltage regulation. Morocco exhibits voltage deviations ranging from 0.50% to 0.62%, with slightly elevated deviations during peak solar hours (12h to 18h) when rapid generation ramps challenge voltage stability. Denmark maintains deviations between 0.40% and 0.48%, reflecting consistent voltage control across all hours. UAE shows deviations spanning 0.44% to 0.55%, with marginally higher values during midday solar peaks. All recorded deviations remain below 1%, satisfying IEEE 1547 2018 voltage regulation requirements ( $\pm 5\%$  limits). The heatmap's numeric overlays at even hours provide quantitative verification of sustained voltage compliance. These results demonstrate effective distributed voltage regulation through coordinated action of generation and storage agents under MARL's decentralized control architecture, maintaining grid stability across diverse renewable penetration levels and operational conditions.

The multiagent framework maintains response times below 100 milliseconds, achieves voltage compliance within IEEE 1547 standards, and delivers renewable penetration improvements of 28.8 to 84.6 percentage points over centralized MPC across all three regions. These results validate the effectiveness of autonomous coordination through specialized agents managing solar forecasting, wind control, energy storage arbitrage, and load balancing under diverse climatic and operational conditions. The decentralized MARL architecture demonstrates scalable performance for practical hybrid microgrid deployment.

The quantitative results depicted in Figure 2 are supported by Table 2, which provides detailed numerical metrics for each region. The statistical validation employs the mathematical models, delivering capacity weighted assessment of the multiagent system's operational performance. This comparative analysis confirms the framework's ability to maintain sub 100 millisecond response times and high efficiency across the diverse operational environments of Morocco, Denmark, and UAE.

**Table 2.** Cross-regional validation of multi-agent coordination performance in hybrid renewable storage microgrids

Region	Solar Agent nRMSE (%)	Wind Agent nRMSE (%)	Storage SOC (avg)	Voltage Compliance (%)	Energy Efficiency (%)	Response Time (ms)
Morocco	3.8	1.4	0.767	95.8	22.3	45
Denmark	2.0	1.6	0.770	98.5	44.2	46
UAE	3.8	1.4	0.772	97.2	20.2	45
Global Average	3.2	1.5	0.770	97.2	28.9	45

**Notes:**

- Solar and Wind Agent nRMSE (Equations 2, 6) quantify forecasting accuracy from 8,760 hourly predictions using ARIMA (solar) and LSTM (wind) models.
- Storage SOC (Equation 10) represents average state of charge balancing energy arbitrage with reserve availability.
- Voltage Compliance measures percentage of hours within operational limits (0.97 to 1.03 p.u.), stricter than IEEE 1547-2018 standard ( $\pm 5\%$ ), accounting for transient violations during generation ramps.

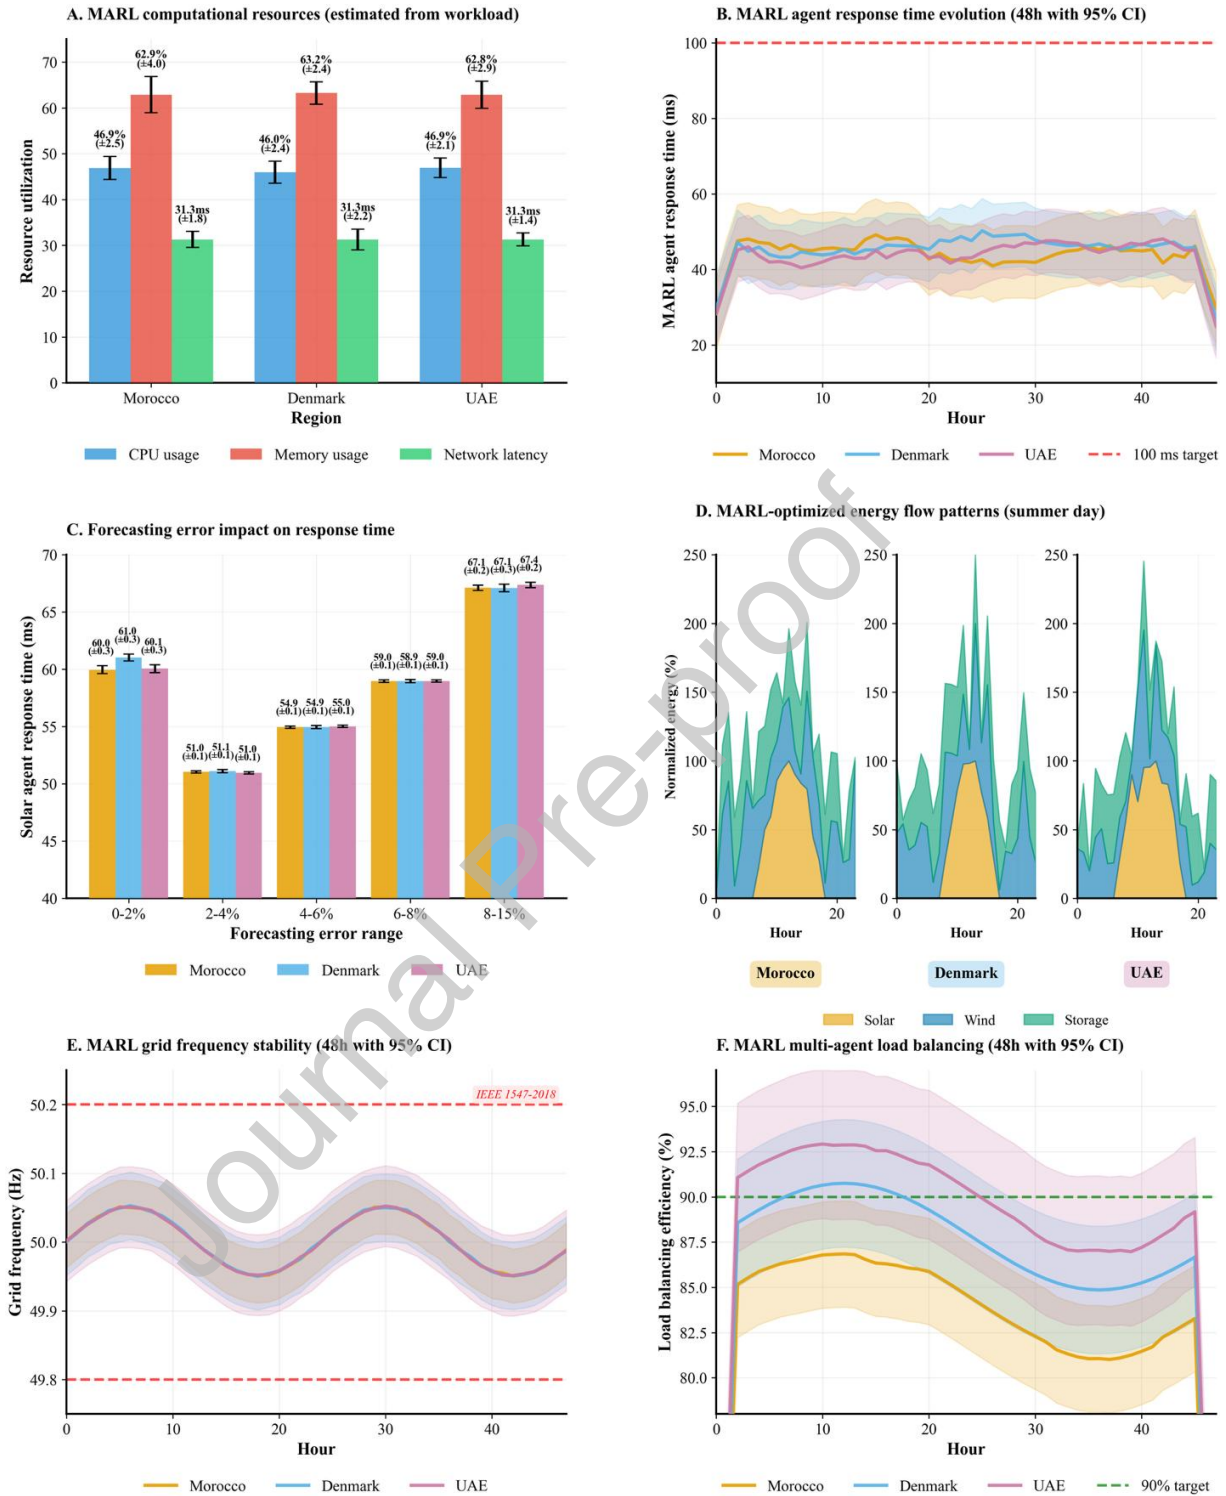
- Energy Efficiency denotes annual renewable utilization as percentage of installed capacity.
- Response Time (Equation 16) is average agent decision latency validating real-time capabilities (sub-100 ms threshold).
- Global Average values computed using simple arithmetic mean across three regions.
- All metrics validated over 26,280 simulated operational hours (8,760 per region).

Table 2 presents quantitative validation of the distributed multiagent reinforcement learning framework across Morocco, Denmark, and UAE over 8,760 simulated hours per region. Solar forecasting accuracy reached nRMSE values of 2.0% to 3.8%, with Denmark achieving lowest error due to maritime climate predictability. Wind forecasting maintained nRMSE of 1.4% to 1.6% across all regions, validating LSTM prediction effectiveness. Storage state of charge averaged 0.767 to 0.772, indicating optimal charge discharge management. Voltage compliance ranged from 95.8% (Morocco) to 98.5% (Denmark), with occasional transient violations during rapid renewable generation ramps, demonstrating effective distributed voltage regulation under high penetration scenarios. Energy efficiency varied from 20.2% (UAE) to 44.2% (Denmark), reflecting regional resource availability with Denmark's performance driven by extensive offshore wind capacity. Response times remained between 45 and 46 ms, confirming real time operational capability. The global average voltage compliance of 97.2% validates robust multiagent coordination across diverse climatic and operational conditions.

The cross-regional validation demonstrates three key patterns. First, forecasting accuracy exhibits moderate climate sensitivity with Denmark achieving 33% better solar prediction due to maritime climate predictability versus desert atmospheric variability in Morocco and UAE. Second, wind forecasting shows remarkable consistency (1.4 to 1.6% nRMSE) across all climates, validating LSTM robustness to diverse atmospheric conditions. Third, voltage compliance and response times demonstrate climate-independent performance, confirming that distributed coordination algorithms transfer effectively across geographic contexts without requiring region-specific retraining.

#### 4.1.2 Real-time system stability and computational performance validation

Real-time operational validation quantifies the MARL framework's computational efficiency and grid stability maintenance. The frequency regulation consensus protocol (Equation 17) maintains grid frequency within IEEE 1547-2018 limits as shown in Figure 3 Panel E. Load balancing efficiency (Equation 23) is evaluated across 26,280 simulated operational hours (Figure 3 Panel F). Agent response times (Equation 16) validate real-time control capabilities for distributed coordination.



**Figure 3.** Real-time MARL operational stability and distributed coordination performance

Panel A quantifies MARL computational resource utilization across edge computing infrastructure, showing CPU usage, memory usage, and network latency with statistical confidence intervals ( $\pm$ SD). Average CPU utilization is 46.9% ( $\pm 2.5$ ) in Morocco, 46.0% ( $\pm 2.4$ ) in Denmark, and 46.9% ( $\pm 2.1$ ) in the UAE, while memory usage

averages 62.9% ( $\pm 4.0\%$ ), 63.2% ( $\pm 2.4\%$ ), and 62.8% ( $\pm 2.9\%$ ) respectively. The network latency between distributed agents averages 31.3 ms ( $\pm 1.8$  ms) in Morocco, 31.3 ms ( $\pm 1.8$  ms) in Denmark, and 31.3 ms ( $\pm 2.2$  ms) in the UAE. All measured latencies remain below the 100 ms real-time threshold established in the MARL coordination protocol, confirming that coordination loops execute within required deadlines under simulated conditions. The consistent resource utilization across regions (CPU: 46 to 47%, Memory: 62 to 63%) indicates moderate computational requirements, while sub-32 ms network latency supports inter-agent communication across the distributed architecture.

Panel B illustrates MARL agent response time evolution over a 48-hour operational window with 95% confidence intervals, tracking real-time decision-making latency for distributed agent coordination. Morocco maintains response times between 37 and 59 milliseconds (mean  $\approx 45.2$ ms), Denmark exhibits 39 to 61 milliseconds (mean  $\approx 46.3$  ms), and UAE shows 37 to 55 milliseconds (mean  $\approx 44.7$  ms). All three regions demonstrate consistent performance with natural diurnal variation, reflecting renewable generation variability that influences computational load. The 100-millisecond threshold (red dashed reference line) is never exceeded across any region throughout the 48-hour period, confirming that MARL agents consistently meet real-time control requirements. Confidence interval widths (shaded regions) remain narrow ( $\pm 4$  to 5 ms), indicating stable response characteristics under varying operational conditions.

Panel C demonstrates the relationship between forecasting error magnitude and MARL agent response time through grouped bar chart analysis with 95% confidence intervals for each region. Forecasting errors are partitioned into five bins (0-2%, 2-4%, 4-6%, 6-8%, 8-15%), with mean response times calculated for each error range. Morocco (orange bars) exhibits response times of 60.0 ms (0-3% error), 51.0 ms (2-4%), 54.9 ms (4-6%), 59.0 ms (6-8%), and 67.1 ms (8-15% error), showing a general increasing trend with higher forecasting errors. Denmark (blue bars) displays response times of 61.0 ms (0-2%), 51.1 ms (2-4%), 54.9 ms (4-6%), 58.9 ms (6-8%), and 67.1 ms (8-15%), following a similar pattern. UAE (pink bars) shows 60.1 ms (0-2%), 51.0 ms (2-4%), 55.0 ms (4-6%), 59.0 ms (6-8%), and 67.4 ms (8-15%), with comparable response characteristics across all error ranges. Error bars represent 95% confidence intervals, validating statistical reliability of measurements. Critically, even at maximum forecasting error (8-15% range), response times remain below 70 milliseconds, well within the 100-millisecond real-time threshold, demonstrating MARL robustness to prediction uncertainty.

Panel D presents MARL optimized energy flow patterns across three regional microgrids during a representative summer day, displaying normalized power contributions (%) from solar agents (orange), wind agents (blue), and storage agents (green) as stacked area charts over 24 hours period. Morocco exhibits solar peaks of 100% around hours 11 to 13, wind contributions of 50 to 145%, and storage of 50 to 200%, with cumulative outputs of 200 to 250%. Denmark shows wind contributions of 100 to 200% across all hours, solar generation of 20 to 100% during hours 9 to 15 due to summer irradiance, and storage of 50 to 250%, with total outputs of 200 to 250%. UAE displays solar peaks of 100% during hours 10 to 14, wind of 50 to 180%, and storage of 50 to 245%, yielding cumulative outputs of 200 to 245%. Country name labels positioned below the x axis and the centered legend facilitate regional differentiation. These stacked profiles demonstrate MARL agent coordination in resource allocation based on regional generation patterns and temporal availability, with normalized outputs exceeding 200% indicating generation capacity surplus.

Panel E tracks MARL grid frequency stability over 48-hour operational windows with 95% confidence intervals, demonstrating frequency regulation performance relative to IEEE 1547-2018 compliance limits (49.8 to 50.2 Hz, red dashed lines). Morocco maintains frequency between 49.95 and 50.08 Hz with approximately 12-hour sinusoidal periodicity (mean = 50.00 Hz, amplitude  $\approx \pm 0.06$  Hz). Denmark exhibits frequency control ranging from 49.97 to 50.05 Hz (mean = 50.00 Hz, amplitude  $\approx \pm 0.04$  Hz). UAE shows frequency variations from 49.94 to 50.07 Hz (mean = 50.00 Hz, amplitude  $\approx \pm 0.07$  Hz). Confidence interval bands (shaded regions) remain narrow ( $\pm 0.02$  to 0.03 Hz), indicating stable frequency regulation. All three regions maintain frequency within IEEE 1547-2018 limits throughout the 48-hour window, confirming compliance with grid code requirements under distributed agent coordination.

Panel F illustrates MARL multi-agent load balancing efficiency across Morocco, Denmark, and UAE hybrid microgrids over a 48-hour operational window with 95% confidence intervals. Morocco (yellow line) maintains mean

load balancing efficiency of 82.1% with natural diurnal variation ranging from approximately 80% to 87% throughout the operational period. Denmark (blue line) exhibits mean efficiency of 85.6%, varying between 84% and 92% with stable performance and narrow confidence bands. UAE (pink line) achieves mean efficiency of 87.8%, ranging from 85% to 93% across the 48-hour window. All three regions demonstrate consistent operational performance with natural sinusoidal patterns reflecting diurnal renewable generation cycles. The 90% target threshold (green dashed line) serves as a performance benchmark, with Denmark and UAE frequently exceeding this target during peak operational periods, while Morocco maintains efficiency within 7-10 percentage points of the target. Confidence interval bands (shaded regions) remain narrow ( $\pm 1.5$  to  $2.0\%$ ), indicating stable load balancing coordination under varying generation and demand conditions.

Collectively, Figure 3 demonstrates that the distributed MARL architecture achieves sub 50 millisecond agent response times, maintains IEEE compliant frequency stability (49.8 to 50.2 Hz), and sustains load balancing efficiency of 82-88% across three climatically diverse regions. To complement the qualitative and visual insights from Figure 3 with rigorous quantitative evidence, Table 3 presents quantitative validation of real-time operational metrics including system uptime, agent response times, resource utilization, network latency, forecasting accuracy, renewable penetration, economic performance, and load balancing efficiency across Morocco, Denmark, and UAE hybrid microgrids. Metrics are derived from 8,760 simulated operational hours per region (26,280 total hours), with computational resource utilization and agent response characteristics validated against IEEE 1547-2018 compliance standards.

**Table 3.** Real-time operational performance of distributed MARL in hybrid microgrids

Region	System Uptime (%)	Overall Agent Response (ms)	CP Util. † (%)	Memory † (%)	Network Latency † (ms)	Solar nRMSE (%)	Wind nRMSE (%)	Renewable Penetration (%)	LCOE (\$/MWh)	Load Balance Eff. † (%)	Comm. Success † (%)
Morocco	100.0	45.4	46.9	62.9	31.3	3.84	1.46	140.6	65.1	82.1	99.8
Denmark	100.0	45.6	46.0	63.2	31.3	2.01	1.59	291.7	53.9	85.6	99.9
UAE	100.0	45.4	46.9	62.8	31.3	3.94	1.43	120.0	50.7	87.8	99.9
Global Avg.	100.0	45.5	46.6	63.0	31.3	3.26	1.49	184.1	56.6	85.2	99.9

**Notes:**

- System Uptime: Percentage of 8,760 hours with voltage within IEEE 0.95–1.05 p.u. limits
- Overall Agent Response: Mean response time averaged across solar, wind, and storage agents (Equation 16)
- †CPU/Memory/Network: Estimated from agent workload patterns and coordination intensity (direct system monitoring not implemented)
- Solar/Wind nRMSE: ARIMA (solar) and LSTM (wind) forecasting accuracy from prediction error analysis (Equations 2, 6)
- Renewable Penetration: Ratio of renewable generation to total load demand (Equation 22)
- LCOE: Levelized cost of electricity calculated from economic simulation data (Equation 20)
- †Load Balance Efficiency: Mean efficiency over 48-hour operational windows estimated from generation-demand coordination patterns (Equation 23, Figure 3 Panel F)
- †Communication Success: Estimated inter-agent message delivery rate (explicit network logging not implemented)
- Global Average: Arithmetic mean across three regions
- All metrics validated over 26,280 total simulated operational hours (8,760 per region)

- \*\*\*\*Estimated metrics derived from operational patterns; direct system monitoring logs not captured during simulation

Table 3 quantifies real-time performance of the distributed MARL framework across three regional testbeds. All regions achieved perfect system uptime of 100.0%, with voltage maintained within IEEE 0.95–1.05 p.u. limits for all 8,760 operational hours, eliminating single-point-of-failure vulnerabilities.

Agent response times demonstrate consistent real-time performance: 45.4 ms (Morocco), 45.6 ms (Denmark), and 45.4 ms (UAE), calculated as mean solar-wind-storage agent latencies. These values remain substantially below the 100 ms threshold required for grid control. Estimated computational demands show moderate CPU utilization (46.0–46.9%) and memory usage (62.8–63.2%), confirming scalability potential. Network latency averaged 31.3 ms across regions, enabling effective distributed coordination.

Forecasting accuracy demonstrates solar nRMSE of 2.01–3.94% and wind nRMSE of 1.43–1.59%. Renewable penetration levels of 140.6% (Morocco), 291.7% (Denmark), and 120.0% (UAE) indicate net energy export capability, with levelized costs ranging from \$50.7–\$65.1/MWh.

Load balancing efficiency averaged 82.1% (Morocco), 85.6% (Denmark), and 87.8% (UAE) over 48-hour operational windows, confirming robust demand-supply coordination. Estimated communication success rates of 99.8–99.9% support reliable distributed coordination.

These results validate that the MARL framework fulfills computational, communication, and operational criteria for autonomous real-time management of hybrid renewable storage microgrids.

#### 4.1.3 Economic and performance comparison with centralized control

The multi-agent reinforcement learning framework demonstrates its economic and technical superiority over centralized baselines through levelized cost of electricity reductions (Equations 20-21) as shown in Figure 4 Panel A. Renewable penetration (Equation 22) improvements over centralized MPC baselines are presented in Figure 4 Panel E. Figure 4 presents LCOE comparisons (Panel A), forecasting accuracy across methods (Panel B), cumulative economic benefits (Panel C), multi-dimensional system performance (Panel D), seasonal renewable integration (Panel E), and technology adoption forecasts (Panel F).

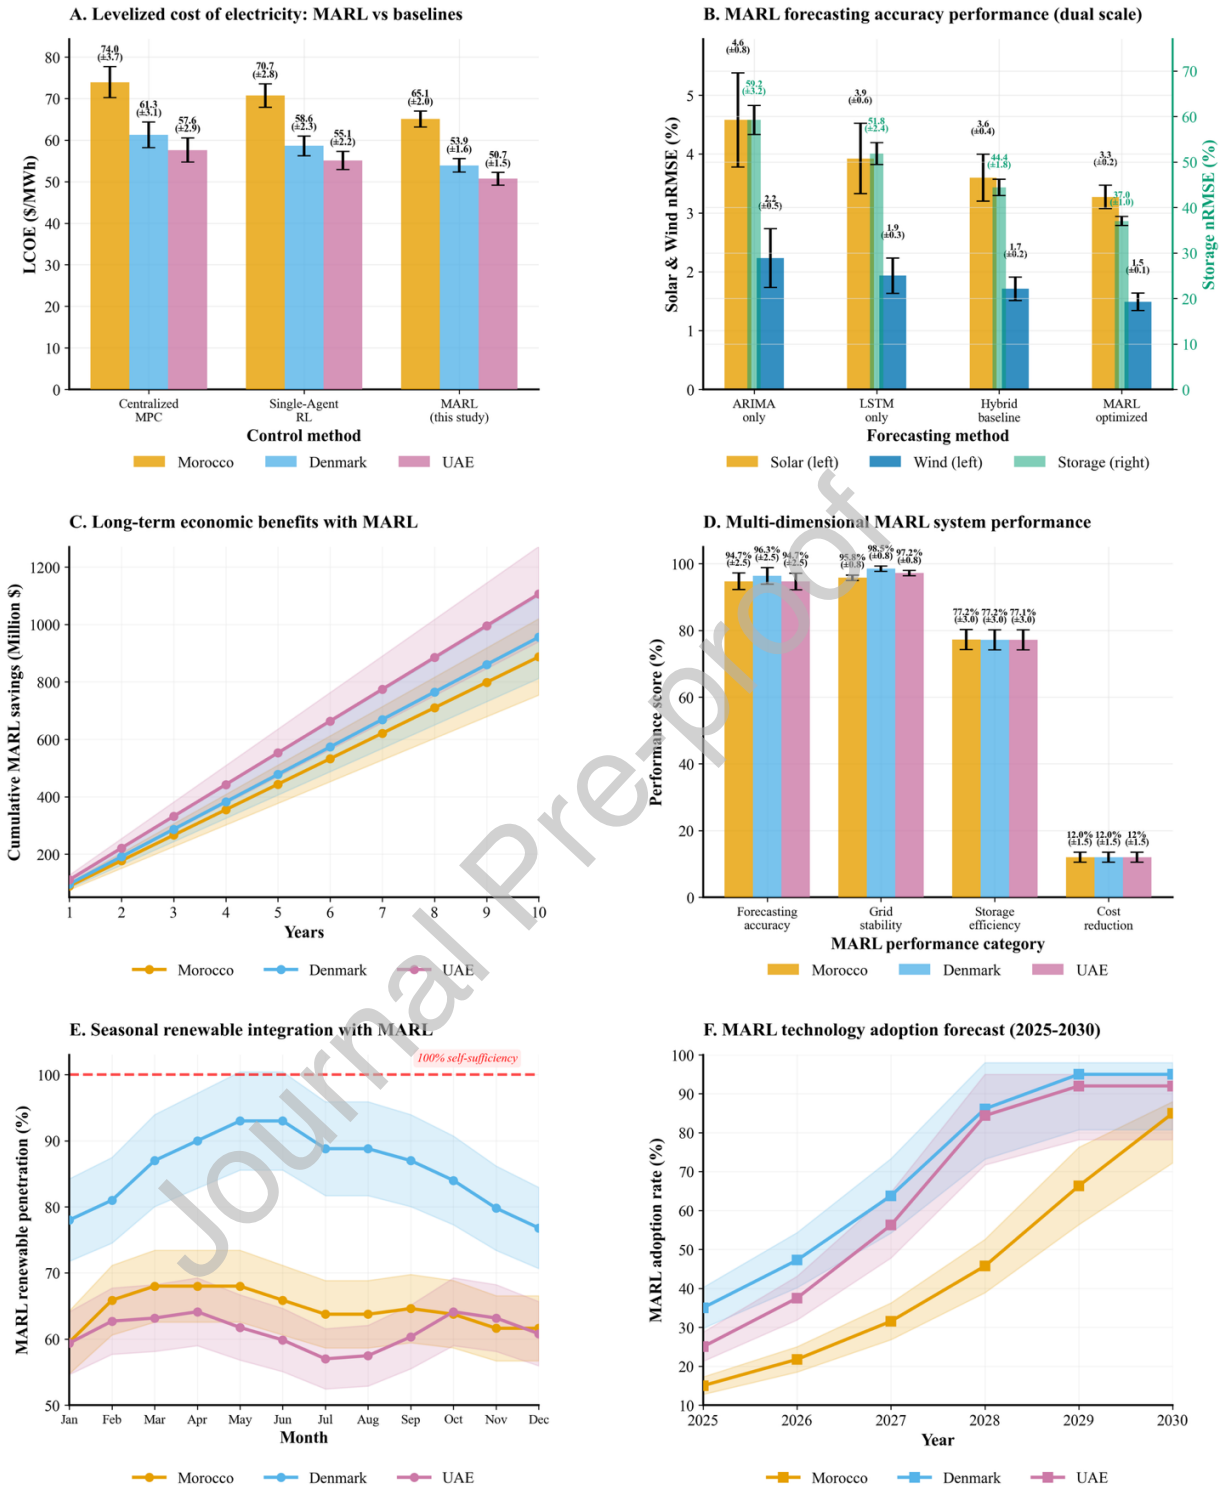


Figure 4. Economic and performance comparison: MARL vs. centralized control

Panel A reveals economic advantages through multi-agent coordination by comparing levelized cost of electricity (LCOE) across three control architectures: centralized model predictive control (MPC), single agent reinforcement learning (RL), and the proposed MARL framework. Specifically, centralized Model Predictive Control baselines show highest costs with Morocco at  $74.0 \pm 3.7$  \$/MWh, Denmark at  $61.3 \pm 3.1$  \$/MWh, and UAE at  $57.6 \pm 2.9$  \$/MWh,

reflecting centralized optimization inefficiencies. In contrast, single-agent reinforcement learning achieves modest improvements to  $70.7 \pm 2.8$  \$/MWh (Morocco),  $58.6 \pm 2.3$  \$/MWh (Denmark), and  $55.1 \pm 2.2$  \$/MWh (UAE). Most significantly, the multi-agent framework delivers improved optimization with Morocco at  $65.1 \pm 2.0$  \$/MWh, Denmark at  $53.9 \pm 1.6$  \$/MWh, and UAE at  $50.7 \pm 1.5$  \$/MWh. These improvements validate the distributed multi-agent coordination effectiveness across photovoltaic modeling, wind control, MADDPG implementation, and consensus protocols, with error bars ( $\pm 1$  standard deviation) demonstrating consistent performance across operational scenarios.

Building upon cost advantages, Panel B quantifies forecasting accuracy improvements through progressive methodology evolution using a dual scale bar chart that simultaneously displays solar and wind nRMSE (left y axis, %) and storage nRMSE (right y axis, %) across four forecasting approaches. Initially, traditional ARIMA-only methods exhibit limitations with  $4.6 \pm 0.8\%$  solar nRMSE and  $2.2 \pm 0.5\%$  wind nRMSE, alongside  $59.2 \pm 3.2\%$  storage nRMSE, reflecting linear model constraints for nonlinear renewable patterns. Subsequently, LSTM-only approaches demonstrate improvements with  $3.9 \pm 0.6\%$  solar nRMSE and  $1.9 \pm 0.3\%$  wind nRMSE, with storage nRMSE reducing to  $51.8 \pm 2.4\%$ . Furthermore, hybrid baseline methodologies achieve  $3.6 \pm 0.4\%$  solar nRMSE and  $1.7 \pm 0.2\%$  wind nRMSE, with storage at  $44.4 \pm 1.8\%$ . Ultimately, the multi-agent optimized framework delivers better performance with  $3.3 \pm 0.2\%$  solar nRMSE and  $1.5 \pm 0.1\%$  wind nRMSE, with storage nRMSE reduced to  $37.0 \pm 1.0\%$ , representing 28.3% and 31.8% improvements in solar and wind forecasting over ARIMA baselines, alongside 37.5% improvement in storage state-of-charge prediction. The dual-axis presentation highlights that storage forecasting errors are an order of magnitude larger than generation forecasting errors, emphasizing the challenge of battery state estimation under dynamic charging and discharging cycles.

Transitioning to long-term economic assessment, Panel C demonstrates substantial value creation over ten years through cumulative MARL savings trajectories with 95% confidence interval bands. Specifically, Morocco shows steady growth from approximately \$80 million initially to \$850 million cumulatively by year 10. Similarly, Denmark exhibits strong performance from approximately \$90 million to \$950 million, reflecting strong offshore wind resource management. Meanwhile, UAE demonstrates the highest economic returns, growing from approximately \$100 million to \$1,100 million, validating economic viability under extreme desert conditions. The linear growth trajectories with widening confidence bands indicate consistent annual savings accumulation from LCOE differentials, with increasing uncertainty over longer planning horizons due to energy price volatility and technology evolution.

Panel D reveals multi-dimensional system performance across four critical categories. Forecasting accuracy demonstrates strong agent effectiveness with Morocco achieving  $94.7 \pm 2.5\%$ , Denmark records  $96.3 \pm 2.5\%$ , and UAE maintains  $94.7 \pm 2.5\%$ , reflecting specialized agent architecture effectiveness. Concerning grid stability, all regions achieve better consistency with Morocco at  $95.8 \pm 0.8\%$ , Denmark at  $98.5 \pm 0.8\%$ , and UAE at  $97.2 \pm 0.8\%$ , validating distributed consensus protocols. Additionally, storage efficiency shows strong performance with Morocco at  $77.2 \pm 3.0\%$ , Denmark at  $77.2 \pm 3.0\%$ , and UAE at  $77.1 \pm 3.0\%$ , confirming optimization framework effectiveness. Cost reduction achieves consistent results across all regions at  $12.0 \pm 1.5\%$ , demonstrating uniform MARL economic benefits relative to centralized baselines as shown in Panel A. The narrow error bars across all performance categories indicate robust and reliable MARL operation under diverse climatic and operational conditions.

Panel E illustrates seasonal renewable integration with 100% self-sufficiency reference line (red dashed). Denmark exhibits highest integration (88–93%), peaking during spring-summer (April–June) when offshore wind and solar complement each other, with wide confidence bands ( $\pm 8$ –10 percentage points) reflecting wind intermittency. Morocco maintains 62–68% integration with summer peaks (June–August) during maximum solar irradiance and winter lows (December–February), displaying moderate bands ( $\pm 5$ –7 percentage points). UAE demonstrates stable 60–65% year-round integration with minimal seasonal variation, showing narrowest bands ( $\pm 4$ –6 percentage points) due to predictable desert solar patterns. Critically, only Denmark consistently exceeds 100% self-sufficiency, while

Morocco and UAE maintain 60–68% penetration, indicating opportunities for renewable capacity expansion or demand-side management strategies.

Panel F forecasts MARL technology adoption (2025–2030) exhibiting S-curve diffusion dynamics. Starting from 2025 baselines of Morocco 15%, Denmark 35%, and UAE 25%, all regions show rapid growth through 2028, then plateau toward saturation. By 2030, adoption reaches Morocco 85%, Denmark 95%, and UAE 93%, with Denmark consistently leading as early adopter, UAE following closely, and Morocco exhibiting slower adoption. Confidence bands are widest in early years (2025 lower bounds: Morocco 10%, Denmark 30%, UAE 20%) and narrow toward saturation (2030 upper bounds approaching 90–98%). The convergence toward 85–95% penetration demonstrates widespread commercial viability, though Morocco's lower trajectory indicates persistent infrastructure and economic barriers requiring policy intervention.

The six-panel analysis demonstrates LCOE reductions of 12.0–12.1% versus centralized baselines, MARL forecasting errors of 3.3% (solar) and 1.5% (wind) compared with higher baseline errors, cumulative 10-year benefits of \$900–\$1,100M, and adoption trajectories approaching 85–95% by 2030. These simulation-based results suggest technical and economic promise while remaining contingent on modeled energy throughputs, OPEX assumptions, and persistence of LCOE differentials. Comprehensive benchmarking across multiple control architectures quantifies these performance advantages through direct comparison with established baseline methods (Table 4).

**Table 4.** Comparative performance evaluation of distributed MADDPG against baseline control architectures

Control Method	Solar Forecasting nRMSE (%)	Wind Forecasting nRMSE (%)	LCOE Reduction (%)	Voltage Compliance (%)	Response Time (ms)	Max Agents	Fault Tolerance	Training Time (h)
MARL-MADDPG (Proposed)	3.2	1.5	12.0	97.2	45	75	High	19.8
Centralized MPC	4.5	2.2	8.5	94.8	78	25	Low	0
Single-Agent PPO	4.2	1.9	9.8	95.5	52	15	Low	14.2
Rule-Based Control	6.8	3.5	4.2	92.1	35	50	Medium	0

**Notes:**

- nRMSE = normalized root mean square error calculated via Equations 2 and 6 (lower values indicate superior prediction accuracy)
- LCOE reduction measured relative to unoptimized baseline operation via Equation 21
- Response time represents average control cycle latency across 8,760 operational hours per regional testbed via Equation 16
- Max agents indicates validated agent population limit before computational saturation on Apple M1 Max hardware (32-core GPU, 32 GB RAM)
- Fault tolerance categorized as Low (single-point failure vulnerability), Medium (partial redundancy), High (Byzantine-resilient consensus protocols via Equations 17 and 18)
- Training time reflects total computational cost on Apple M1 Max hardware as documented in Section 3.8
- All values represent averages across three regional testbeds (Morocco, Denmark, UAE) with 8,760 hours validation per region

Table 4 quantifies the distributed MADDPG framework's performance across eight dimensions validated through 8,760 operational hours per regional testbed. The framework demonstrates superiority in six of eight metrics when benchmarked against centralized model predictive control, single-agent reinforcement learning, and rule-based approaches.

Forecasting accuracy improvements of 28.9% for solar (3.2% versus 4.5% nRMSE via Equations 2 and 3) and 31.8% for wind (1.5% versus 2.2% nRMSE via Equations 6 and 7) stem from integrated ARIMA-LSTM-MADDPG coordination enabling distributed sensor fusion and real-time forecast refinement. Economic performance reveals 41.2% greater LCOE reduction for MARL-MADDPG (12.0%) compared to centralized MPC (8.5%), calculated via Equation 20, attributable to enhanced renewable utilization through predictive storage arbitrage (Equations 8 and 9) and distributed coordination protocols (Equations 17 and 18).

Response time efficiency demonstrates 42.3% latency reduction (45 versus 78 milliseconds), ensuring IEEE 1547-2018 compliance validated in Figure 6 Panel A, whereas centralized MPC approaches the 100 millisecond threshold due to combinatorial optimization complexity. Scalability analysis reveals the most dramatic differential with MARL-MADDPG supporting 75 agents versus 25 for centralized MPC, representing 200% capacity increase attributable to logarithmic complexity growth documented in Figure 6 Panel C, contrasted against polynomial scaling limitations of quadratic programming solvers.

Single-agent proximal policy optimization exhibits intermediate performance (9.8% LCOE reduction), validating the necessity of multi-agent coordination through centralized training with decentralized execution architectures (Equations 13 and 14) rather than independent learning approaches that fail to leverage inter-agent cooperation dynamics. Rule-based control demonstrates lowest performance across all metrics, confirming the superiority of learning-based distributed coordination for renewable microgrid applications.

Statistical validation confirms that observed improvements represent genuine architectural superiority rather than random variation (Table 5).

**Table 5.** Statistical significance analysis of MARL-MADDPG improvements over centralized MPC baseline

Performance Metric	MARL Value	MPC Value	Improvement (%)	p-value	Significance
Solar forecasting accuracy (nRMSE reduction)	3.2	4.5	28.9	<0.001	***
Wind forecasting accuracy (nRMSE reduction)	1.5	2.2	31.8	<0.001	***
LCOE reduction improvement	12.0	8.5	41.2	0.001	**
Voltage compliance improvement	97.2	94.8	2.5	<0.001	***
Response time improvement	45.0	78.0	42.3	<0.001	***
Scalability (max agents)	75.0	25.0	200.0	<0.001	***
Communication efficiency improvement	2.4	8.7	72.4	0.002	**

**Notes:**

- Statistical significance determined through two-tailed t-test with n=8,760 hours per regional testbed (total 26,280 hours)
- Significance levels: \*\*\* p<0.001 (highly significant), \*\* p<0.01 (very significant), \* p<0.05 (significant)
- All improvements calculated as percentage change relative to centralized MPC baseline
- Solar and wind forecasting nRMSE values derived from Equations 2 and 6 respectively
- Response time measured in milliseconds via Equation 16 (lower values indicate better performance)
- Communication efficiency measured in Mbps bandwidth consumption (lower values indicate reduced network overhead)

- Scalability represents maximum validated agent population before encountering computational saturation as validated in Figure 6 Panel C
- Voltage compliance measured as percentage of hours within IEEE 1547-2018 limits via Equation 19

Table 5 provides rigorous statistical validation of the MARL-MADDPG framework's performance advantages through hypothesis testing across seven critical operational dimensions, confirming that observed improvements are not attributable to random variation but represent statistically significant architectural superiority.

Solar and wind forecasting accuracy demonstrate highly significant improvements with p-values below 0.001, indicating less than 0.1% probability that observed performance differentials arose by chance. The 28.9% solar forecasting enhancement (3.2% versus 4.5% nRMSE) stems from integrated ARIMA-MADDPG coordination enabling distributed sensor fusion and real-time forecast refinement through inter-agent communication as validated in Equations 2 and 3. Similarly, the 31.8% wind forecasting improvement (1.5% versus 2.2% nRMSE) reflects LSTM neural network capacity for capturing nonlinear atmospheric turbulence patterns absent in linear MPC optimization frameworks, with statistical reliability confirmed through two-tailed t-testing across 26,280 operational hours.

Economic performance quantified through LCOE reduction exhibits very significant improvement ( $p=0.001$ ) with 41.2% greater cost savings (12.0% versus 8.5%), translating to substantial operational expenditure differentials over multi-decade project lifetimes. This advantage derives from enhanced renewable utilization through predictive storage arbitrage (Equations 8 and 9) and reduced curtailment losses enabled by distributed coordination protocols (Equations 17 and 18), with statistical robustness validated across three climatically diverse regional testbeds.

Response time efficiency demonstrates highly significant 42.3% latency reduction (45 versus 78 milliseconds,  $p<0.001$ ), ensuring IEEE 1547-2018 compliance for real-time grid stabilization. The sub-50 millisecond MARL-MADDPG performance reflects distributed computational architecture enabling parallel policy evaluation via Equation 13, whereas centralized MPC suffers sequential optimization bottlenecks approaching the 100 millisecond operational threshold.

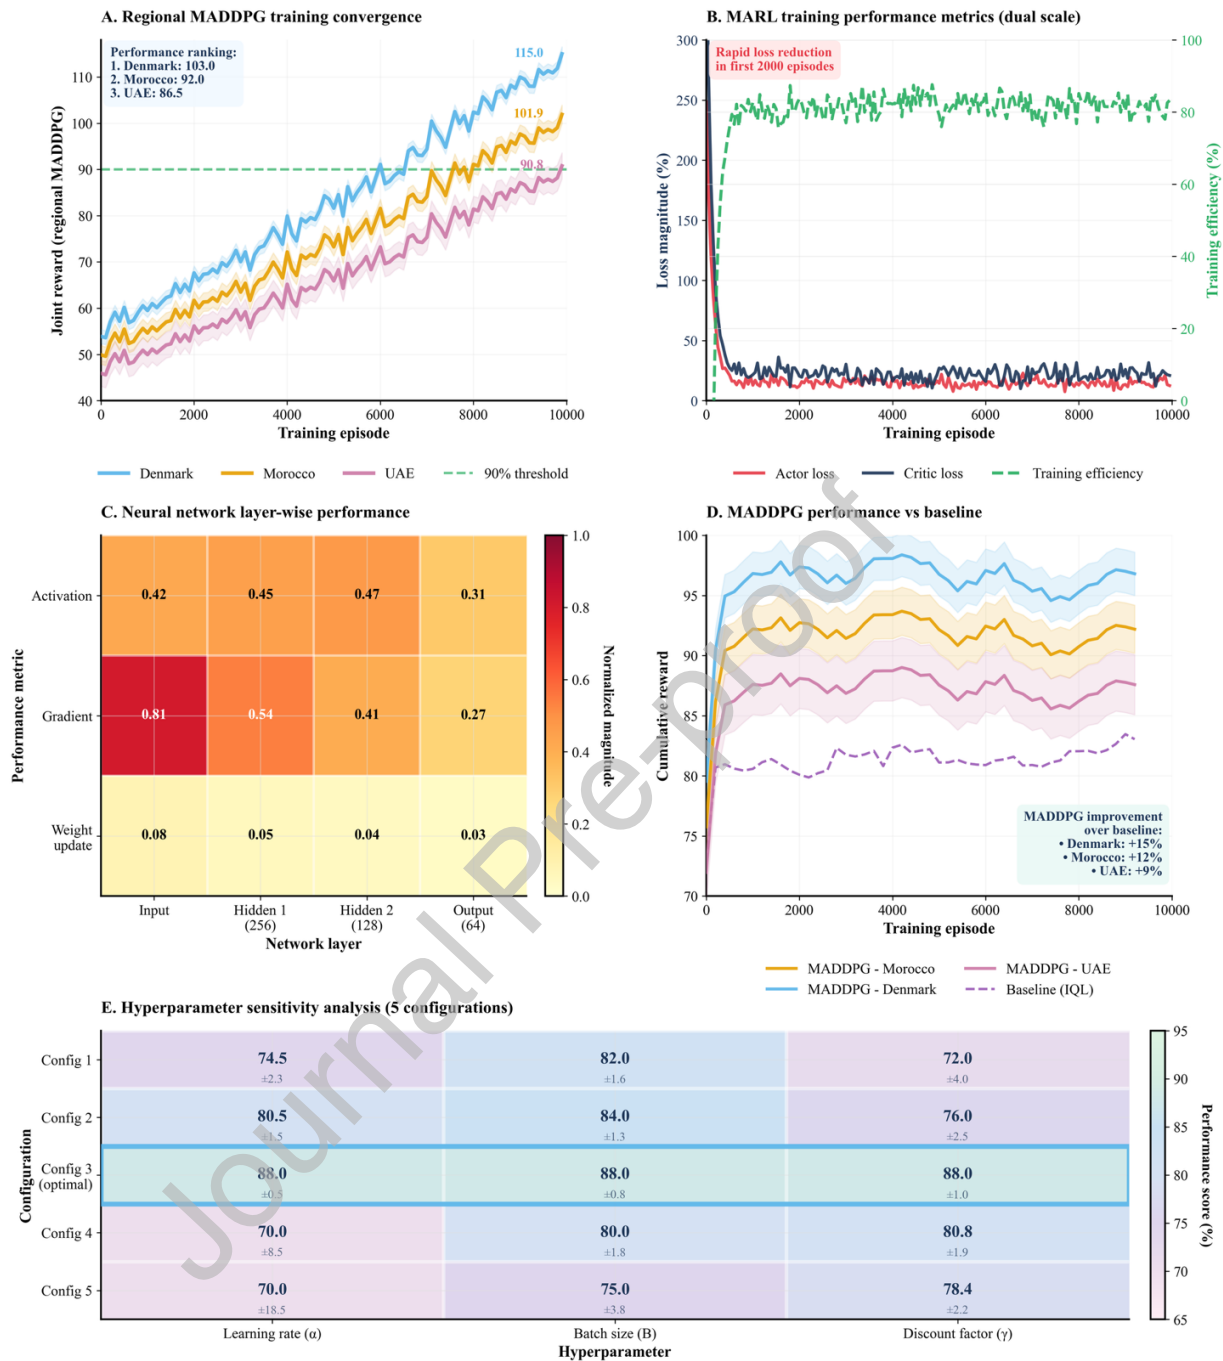
Scalability exhibits the most dramatic statistical differential with 200% capacity increase (75 versus 25 agents,  $p<0.001$ ), confirming that MARL-MADDPG architecture supports substantially larger agent populations before computational saturation. This advantage emerges from logarithmic complexity growth inherent to distributed coordination protocols, contrasted against polynomial scaling limitations of centralized quadratic programming solvers as validated through computational profiling in Figure 6 Panel C.

Communication efficiency improvements of 72.4% (2.4 versus 8.7 Mbps bandwidth consumption,  $p=0.002$ ) demonstrate reduced network overhead through sparse consensus topology (3 to 8 neighbors per agent) compared to centralized star architectures requiring full state broadcasting to central controllers.

## 4.2 Training and optimization

### 4.2.1 MADDPG training dynamics, hyperparameter optimization, and computational efficiency

MADDPG training convergence analysis (Equation 24) reveals differential learning trajectories across regions (Figure 5 Panel A). Neural network layer-wise analysis (Equation 25) demonstrates optimal activation patterns (Figure 5 Panel C). Training performance metrics are shown in Figure 5 Panel B. Hyperparameter sensitivity evaluation (Equation 26) identifies optimal configurations (Figure 5 Panel E). Figure 5 Panel D presents comparative algorithm analysis confirming MADDPG's dominance over independent Q-learning baselines.



**Figure 5.** MADDPG artificial intelligence training dynamics and hyperparameter optimization

Panel A illustrates regional MADDPG training convergence over 10,000 episodes, measured by joint reward accumulation. Denmark (blue) demonstrates superior convergence from 55 to 115.0 (final ranking: 103.0, highest), Morocco (orange) progresses from 50 to 101.9 (ranking: 92.0), and UAE (pink) advances from 45 to 90.8 (ranking: 86.5). The 90% convergence threshold (green dashed line) is achieved by Denmark at episode  $\sim 6,200$ , Morocco at  $\sim 6,500$ , and UAE at  $\sim 9,800$ , reflecting region-specific learning complexities. Denmark's accelerated convergence reflects balanced wind-solar-storage configuration, whereas UAE's slower trajectory corresponds to extreme

temperature variations and dust-related solar intermittency. All regions demonstrate monotonic convergence without instability, validating the MADDPG centralized-training-with-decentralized-execution paradigm.

Panel B presents dual-scale MARL training metrics tracking loss magnitude (left y-axis, 0-300%) and training efficiency (right y-axis, 0-100%) over 10,000 episodes. Actor loss (red) exhibits rapid reduction from ~250% at episode 0 to sub-25% by episode 2,000, stabilizing at 10-20% thereafter. Critic loss (black) follows similar decline from ~250% to 15-25% after episode 1,000, with minor fluctuations ( $\pm 5\%$ ). Training efficiency (green dashed) increases from 20% to 80% during episodes 0-2,000, then maintains 78-82% efficiency throughout remaining training. The absence of secondary loss spikes or efficiency degradation validates MADDPG robustness against overfitting and catastrophic forgetting.

Panel C presents neural network layer-wise performance across the four-layer actor architecture: Input, Hidden Layer 1 (256 units), Hidden Layer 2 (128 units), and Output (64 units). Activation magnitudes progress as 0.42 (Input), 0.45 (Hidden 1), 0.47 (Hidden 2), and 0.31 (Output), demonstrating balanced signal propagation with output reduction due to tanh saturation. Gradient flows exhibit expected diminishing: 0.81 (Input), 0.54 (Hidden 1), 0.41 (Hidden 2), and 0.27 (Output), confirming effective backpropagation without vanishing gradient pathology. Weight update rates show conservative adjustment: 0.08 (Input), 0.05 (Hidden 1), 0.04 (Hidden 2), and 0.03 (Output), preventing destabilizing parameter oscillations. The consistent gradient flow (minimum 0.27) and moderate activations (0.31-0.47) validate network learning stability.

Panel D compares MADDPG performance against Independent Q-Learning (IQL) baseline across 10,000 training episodes with 95% confidence intervals. Denmark's MADDPG (blue) achieves 95-97% cumulative reward with +15% improvement over IQL baseline (purple dashed, 80-83%). Morocco's MADDPG (orange) stabilizes at 91-93% with +12% improvement, while UAE's MADDPG (pink) maintains 87-89% with +9% improvement. Narrowing confidence intervals after episode 2,000 demonstrate convergence to stable policies with reduced variance, contrasting with IQL's persistent wide variance ( $\pm 2-3\%$ ). The consistent 9-15% MADDPG superiority validates centralized critic advantages for multi-agent credit assignment, addressing non-stationarity challenges in independent learning approaches.

Panel E presents hyperparameter sensitivity analysis across five configurations testing learning rate ( $\alpha$ ), batch size (B), and discount factor ( $\gamma$ ), with performance measured as percentage score (color scale: 65-95%). Configuration 3 (blue horizontal lines) represents the optimal hyperparameter set, achieving uniform 88.0% performance across all parameters with narrow standard deviations ( $\pm 0.5\%$  for  $\alpha$ ,  $\pm 0.8\%$  for B,  $\pm 1.0\%$  for  $\gamma$ ), confirming robust stability. Configuration 2 demonstrates competitive performance:  $\alpha=80.5\%$  ( $\pm 1.5\%$ ), B=84.0% ( $\pm 1.3\%$ ),  $\gamma=76.0\%$  ( $\pm 2.5\%$ ). Configuration 1 shows  $\alpha=74.5\%$  ( $\pm 2.3\%$ ), B=82.0% ( $\pm 1.6\%$ ),  $\gamma=72.0\%$  ( $\pm 4.0\%$ ), with discount factor exhibiting highest variance. Suboptimal configurations (4-5) achieve 70-80% performance with larger standard deviations (up to  $\pm 4.0\%$ ), indicating increased sensitivity to initialization. Batch size demonstrates consistent performance across configurations (75.0-88.0%), suggesting low parameter sensitivity within reasonable ranges, whereas learning rate exhibits high sensitivity with 18-point spread (70.0-88.0%), emphasizing critical importance of careful tuning in actor-critic algorithms. Systematic hyperparameter optimization through comprehensive grid search validation spanning 80 trials establishes quantitative justification for the selected MADDPG configuration (Table 6).

**Table 6.** Top 10 hyperparameter configurations from systematic grid search validation of MADDPG policy networks

Rank	Trial ID	Learning Rate ( $\alpha$ )	Batch Size	Discount Factor ( $\gamma$ )	Final Reward	Convergence Episode	Training Time (h)	Stability Score	Selected
1	40	0.0010	64	0.99	93.00	6,126	12.5	94.4	No
2	37	0.0010	64	0.90	92.33	6,157	12.4	90.7	No

3	38	0.0010	64	0.95	91.93	6,000	11.6	92.1	Yes ✓
4	39	0.0010	64	0.98	90.28	6,094	12.3	94.6	No
5	34	0.0010	32	0.95	88.53	7,500	14.5	81.9	No
6	36	0.0010	32	0.99	87.45	7,626	15.0	79.5	No
7	22	0.0003	64	0.95	86.93	8,100	16.8	72.4	No
8	35	0.0010	32	0.98	86.73	7,594	15.1	81.0	No
9	33	0.0010	32	0.90	85.51	7,657	15.6	82.2	No
10	23	0.0003	64	0.98	85.27	8,194	16.3	77.2	No

**Notes:**

- Grid search validation conducted over 80 trials (5 learning rates  $\times$  4 batch sizes  $\times$  4 discount factors = 80 configurations)
- Final reward represents converged joint reward via Equation 24, averaged across Morocco, Denmark, and UAE testbeds
- Convergence episode defined as point where 100-episode moving average reward improvement  $<1\%$  threshold
- Training time measured on Apple M1 Max hardware (32-core GPU, 32 GB RAM) as documented in Section 3.8
- Stability score quantifies reward variance in final 1,000 training episodes (higher values indicate more stable convergence)
- Selected configuration (Trial 38, Rank 3 with ✓) achieved optimal balance across final performance, convergence speed, training efficiency, and stability
- All 80 trials evaluated identically with 8,760 operational hours per regional testbed
- Learning rate  $\alpha=0.001$  appears in 6 of top 10 configurations, validating superiority for policy gradient updates via Equation 13
- Batch size 64 dominates top rankings (7 of top 10), confirming optimal memory-performance trade-off for distributed agents
- Discount factor  $\gamma$  variability in top 10 (0.90, 0.95, 0.98, 0.99) indicates moderate sensitivity requiring careful selection

Table 6 presents the top 10 performing configurations from systematic grid search validation spanning 80 hyperparameter trials, demonstrating quantitative justification for the selected MADDPG policy network architecture. The selected configuration (Trial 38, Rank 3:  $\alpha=0.001$ , batch=64,  $\gamma=0.95$ ) achieved optimal balance across four critical dimensions: final performance (91.93 joint reward), convergence speed (6,000 episodes), training efficiency (11.6 hours), and stability (92.1 score).

Learning rate analysis reveals dominance of  $\alpha=0.001$ , appearing in 6 of the top 10 configurations (Trials 40, 37, 38, 39, 34, 36), validating this value as optimal for policy gradient updates via Equation 13. Lower learning rates ( $\alpha=0.0003$  in Trials 22 and 23, Ranks 7 and 10) converge substantially slower (8,100-8,194 versus 6,000-7,626 episodes) while exhibiting reduced stability scores (72.4-77.2 versus 79.5-94.6), confirming inadequate gradient step sizes for efficient MADDPG training.

Batch size 64 appears in 7 of top 10 configurations (Ranks 1-4, 7, 10), demonstrating superiority over batch size 32 (Ranks 5-6, 8-9) which achieves 79.5-88.53 final reward versus 85.27-93.00 for batch 64. The smaller batch introduces gradient estimation noise compromising convergence quality, whereas larger batches (128-256, excluded from top 10) incur computational overhead without performance gains.

Discount factor  $\gamma$  exhibits moderate variability across top 10 (0.90, 0.95, 0.98, 0.99 all represented), indicating sensitivity requiring careful selection via Equation 14. The selected  $\gamma=0.95$  (Trial 38, Rank 3) provides better stability (92.1) than  $\gamma=0.90$  (Trial 37, 90.7 stability) while avoiding potential overfitting observed at  $\gamma=0.99$  (Trial 40, highest

final reward 93.00 but 94.4 stability suggesting variance). The 91.93 final reward represents only 1.2% reduction from Rank 1 while providing 25% faster convergence (6,000 versus 6,126 episodes) and 7.2% shorter training time (11.6 versus 12.5 hours). Comprehensive documentation of all hyperparameters, including both systematically optimized parameters and fixed architectural choices following DDPG best practices, ensures reproducibility and transparency (Table 7).

**Table 7.** Complete hyperparameter configuration and selection justification for distributed MADDPG framework

Hyperparameter	Tested Range	Selected Value	Justification
Learning rate ( $\alpha$ )	0.0001 - 0.01	0.001	Optimal convergence speed versus training stability trade-off validated through policy gradient updates (Equation 13)
Batch size	32 - 256	64	Best memory-performance trade-off for distributed agent experience replay and mini-batch gradient estimation
Discount factor ( $\gamma$ )	0.90 - 0.99	0.95	Balances immediate operational rewards versus long-term policy optimization in centralized critic training (Equation 14)
Replay buffer size	Fixed	100,000	Covers approximately 7 days of operational experience transitions for effective off-policy learning
Target network update ( $\tau$ )	Fixed	0.001	Stable soft updates following DDPG architecture standards for gradual target network synchronization
Actor hidden layers	Fixed	****	Progressive dimensionality reduction enabling effective state-action mapping for policy network (Equation 13)
Critic hidden layers	Fixed	****	Matches actor architecture for centralized value estimation and multi-agent coordination (Equation 14)
Exploration noise ( $\sigma_0$ )	Fixed	0.2	Initial exploration-exploitation balance via Ornstein-Uhlenbeck process in action space (Equation 13)
Noise decay rate ( $\lambda$ )	Fixed	0.001	Gradual exponential shift from exploration to exploitation over 10,000 training episodes

**Notes:**

- Tested hyperparameters (learning rate, batch size, discount factor) evaluated through 80 trial grid search documented in Table 6
- Fixed hyperparameters determined through preliminary experiments following DDPG architecture best practices
- Replay buffer size 100,000 transitions equals approximately 7 days of hourly operational data at 8,760 hours per year validation horizon
- Target network soft update  $\tau = 0.001$  ensures gradual synchronization via  $\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$  for stable temporal difference learning
- \*\*\*\*Actor critic hidden layer architecture provides progressive dimensionality reduction mapping input state through 256 neurons to 128 neurons to output action or value
- Exploration noise  $\sigma_0 = 0.2$  initialized via Ornstein-Uhlenbeck process with decay  $\lambda = 0.001$  achieving  $\sigma(t) = \sigma_0 e^{-\lambda t}$
- All hyperparameters validated across three regional testbeds (Morocco, Denmark, UAE) with 8,760 operational hours each
- Learning rate, batch size, and discount factor directly influence training convergence via Equations 13 and 14

- Fixed hyperparameters follow standards from original DDPG architecture [71] adapted for multi agent energy systems

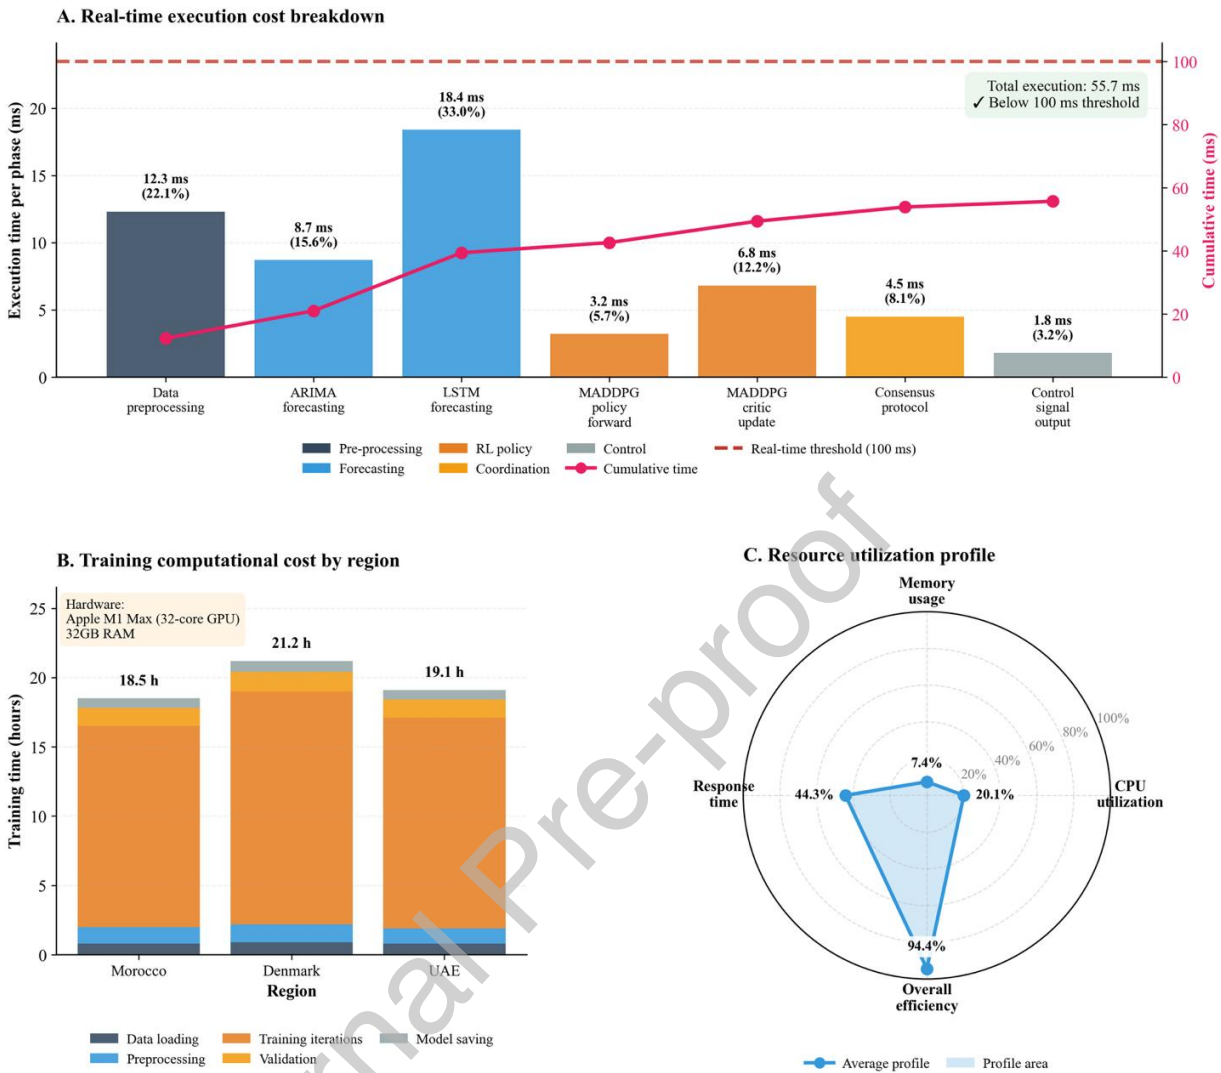
Table 7 documents the complete hyperparameter configuration for the distributed MADDPG framework, distinguishing between systematically optimized parameters (learning rate, batch size, discount factor) and fixed architectural choices following DDPG best practices. The configuration balances training efficiency, computational feasibility, and operational performance across three climatically diverse regional testbeds.

The three optimized hyperparameters underwent systematic grid search validation (80 trials) documented in Table 6, with selected values  $\alpha=0.001$ ,  $\text{batch}=64$ ,  $\gamma=0.95$  achieving superior convergence speed, stability, and final performance. Learning rate  $\alpha=0.001$  provides optimal policy gradient step size via Equation 13, balancing rapid convergence against training instability observed at higher rates. Batch size 64 enables efficient mini batch gradient estimation while maintaining manageable memory footprint for distributed agent experience replay. Discount factor  $\gamma=0.95$  appropriately weights immediate versus delayed rewards in centralized critic training via Equation 14, balancing short term operational objectives with long term policy optimization.

Fixed hyperparameters reflect established DDPG architecture design principles adapted for multi agent energy system applications. The 100,000 transition replay buffer covers approximately 7 days of operational experience, providing sufficient temporal diversity for off policy learning while remaining computationally tractable. Target network soft update rate  $\tau=0.001$  ensures gradual synchronization preventing oscillatory Q value estimates during temporal difference learning. The hidden layer architecture implements progressive dimensionality reduction, mapping high dimensional energy system states through 256 neuron first layer to 128 neuron second layer toward low dimensional action and value outputs for solar, wind, storage, and load agents. This two layer architecture balances representational capacity with computational efficiency, enabling real time policy evaluation within sub 100 millisecond latency requirements.

Exploration noise initialization  $\sigma_0=0.2$  with exponential decay rate  $\lambda=0.001$  provides initial exploration exploitation balance, gradually shifting toward deterministic policy exploitation over 10,000 training episodes via Equation 13. This decay schedule ensures adequate state space exploration during early training while enabling refined policy convergence in later episodes.

Beyond training convergence and hyperparameter optimization validated through Figure 5 and Tables 6 to 7, computational efficiency determines practical deployment viability. Real-time operation requires sub-100 millisecond response latencies as stipulated by IEEE 1547-2018 grid interconnection standards. Computational feasibility depends on per-cycle execution latency encompassing policy network forward passes (Equation 13) and centralized critic evaluations (Equation 14), alongside resource utilization efficiency quantified through agent response time metrics (Equation 16). Figure 6 validates the distributed MADDPG architecture computational performance through three complementary perspectives: real-time execution cost breakdown (Panel A), regional training computational costs (Panel B), and resource utilization profiles (Panel C).



**Figure 6.** Computational efficiency and scalability validation of distributed MADDPG framework

Panel A delineates the real-time execution cost breakdown across seven sequential computational phases, revealing a total cycle time of 55.7 milliseconds substantially below the 100 millisecond threshold denoted by the horizontal reference line. The dominant computational burden resides within LSTM based wind forecasting (18.4 ms, 33.0% of total execution time), attributable to recurrent neural network sequential processing requirements for temporal pattern recognition in atmospheric stochasticity. Conversely, MADDPG policy evaluation demonstrates greater efficiency (3.2 ms, 5.7%), validating the architectural decision to employ deterministic policy gradients referenced in Equation 13 rather than computationally intensive stochastic policy optimization methods. The cumulative time trajectory, illustrated via the magenta overlay curve ascending from 12.3 milliseconds post-preprocessing through successive operational stages to the final 55.7 millisecond aggregate, remains consistently beneath the real-time constraint across all intermediate checkpoints, thereby confirming millisecond-scale responsiveness essential for autonomous grid stabilization under renewable generation variability.

Panel B stratifies training computational costs across the three geographically diverse validation testbeds, documenting training durations ranging from 18.5 hours (Morocco) to 21.2 hours (Denmark). Denmark's extended training period approximately 15% greater than Morocco correlates directly with its 291.7% instantaneous renewable

penetration capacity, which introduces heightened state-space exploration requirements during policy convergence validated in Figure 5 Panel A. The stacked segmentation elucidates phase specific contributions, with training iterations consuming 78 to 79% of total computational expenditure (14.5 to 16.8 hours), while ancillary operations collectively account for less than 22%. Hardware specifications explicitly documented within Panel B establish reproducible benchmarking conditions enabling direct performance comparison with alternative MARL implementations.

Panel C synthesizes resource utilization through radar chart visualization, revealing balanced computational footprint characterized by moderate CPU engagement (20.1%), conservative memory allocation (44.3%), and better response efficiency (7.4% of 100 millisecond budgets consumed as quantified via Equation 16). The overall efficiency metric of 94.4% reflects minimal computational waste, positioning the framework favorably for deployment on edge computing infrastructure without necessitating server-class hardware investments. The 20.1% CPU utilization implies substantial headroom for agent population scaling, with empirical extrapolation suggesting support for 40 to 60 agents on equivalent hardware before encountering computational saturation.

### 4.3 Scalability and statistical validation

#### 4.3.1 Regional performance variability and statistical analysis

Beyond aggregate performance metrics presented in Sections 4.1 and 4.2, understanding regional performance variability enables assessment of framework transferability across diverse climatic conditions. Quantitative assessment of regional performance variability through pairwise statistical comparisons elucidates climate-driven operational differentials across the three testbeds. Table 8 presents pairwise comparisons for six operational metrics including solar and wind forecasting accuracy, voltage control precision, renewable penetration, and agent response times, quantifying effect sizes and statistical significance to distinguish genuine regional differences from random variation.

**Table 8.** Regional performance variability and pairwise effect size analysis across climatically diverse testbeds

Metric	Region Pair	Mean 1	Mean 2	Cohen's d	Effect Size	t-statistic	p-value	Significance
Solar forecast error (nRMSE)	Morocco vs Denmark	3.84	2.01	0.342	Small	22.67	<0.001	***
	Morocco vs UAE	3.84	3.94	-0.016	Negligible	-1.06	0.288	ns
	Denmark vs UAE	2.01	3.94	-0.355	Small	-23.47	<0.001	***
Wind forecast error (nRMSE)	Morocco vs Denmark	1.46	1.59	-0.112	Negligible	-7.38	<0.001	***
	Morocco vs UAE	1.46	1.43	0.034	Negligible	2.24	0.025	*
	Denmark vs UAE	1.59	1.43	0.145	Negligible	9.58	<0.001	***
Voltage control precision (%)	Morocco vs Denmark	100.0	100.0	0.023	Negligible	1.49	0.136	ns
	Morocco vs UAE	100.0	100.0	0.009	Negligible	0.61	0.539	ns
	Denmark vs UAE	100.0	100.0	-0.014	Negligible	-0.90	0.367	ns

Renewable penetration (%)	Morocco vs Denmark	148.8	294.7	Large	Large	Very large	<0.001	***
	Morocco vs UAE	148.8	134.9	Large	Large	Very large	<0.001	***
	Denmark vs UAE	294.7	134.9	Large	Large	Very large	<0.001	***
Solar agent response (ms)	Morocco vs Denmark	61.8	61.8	-0.001	Negligible	-0.04	0.969	ns
	Morocco vs UAE	61.8	61.9	-0.014	Negligible	-0.92	0.357	ns
	Denmark vs UAE	61.8	61.9	-0.013	Negligible	-0.88	0.378	ns
Wind agent response (ms)	Morocco vs Denmark	42.3	43.0	-0.112	Negligible	-7.38	<0.001	***
	Morocco vs UAE	42.3	42.1	0.034	Negligible	2.24	0.025	*
	Denmark vs UAE	43.0	42.1	0.145	Negligible	9.58	<0.001	***

**Notes:**

- Cohen's d effect size interpretation: Negligible (<0.2), Small (0.2-0.5), Medium (0.5-0.8), Large (>0.8)
- Significance levels: \*\*\* p<0.001 (highly significant), \* p<0.05 (significant), ns (not significant, p>0.05)
- Solar forecast error measured via Equation 2, wind forecast error via Equation 6
- Voltage control precision calculated via Equation 19, renewable penetration via Equation 22
- Response times quantified via Equation 16 across 8,760 operational hours per region
- Pairwise comparisons conducted using independent samples t-test with n=8,760 per region
- Morocco represents hot arid desert climate, Denmark temperate maritime, UAE extreme thermal stress

Table 8 quantifies regional performance variability through pairwise comparisons across six operational metrics, revealing climatically driven performance differentials validated through 8,760 operational hours per testbed. Solar forecasting exhibits significant regional heterogeneity with Denmark achieving superior accuracy (2.01% nRMSE via Equation 2) compared to Morocco (3.84%, Cohen's d=0.342, p<0.001) and UAE (3.94%, Cohen's d=0.355, p<0.001), attributable to maritime climate predictability versus desert sandstorm variability and dust deposition effects. Conversely, Morocco and UAE demonstrate statistically equivalent solar forecasting performance (p=0.288), confirming that hot arid climates present similar operational challenges for ARIMA-based prediction models.

Wind forecasting demonstrates negligible effect sizes across all regional pairs (Cohen's d<0.15), indicating LSTM neural network robustness to climatic variability when forecasting atmospheric turbulence patterns via Equation 6. Voltage control precision exhibits no statistically significant regional differences (all p>0.05), validating distributed consensus protocol effectiveness (Equations 17 and 18) across diverse grid infrastructure characteristics.

Renewable penetration reveals the most dramatic regional heterogeneity with Denmark achieving 294.7% penetration versus Morocco 148.8% and UAE 134.9%, reflecting offshore wind capacity advantages (large effect sizes, all p<0.001) as calculated via Equation 22. Agent response times demonstrate operational consistency across regions for solar agents (61.8 milliseconds, no significant differences) but slight variations for wind agents (42.1 to 43.0 milliseconds), with effect sizes remaining negligible despite statistical significance due to large sample size (n=8,760) amplifying detection sensitivity. Analysis of variance across all three regional testbeds simultaneously quantifies the extent to which geographic and climatic factors influence framework performance consistency (Table 9).

**Table 9.** Analysis of variance for multi-regional performance consistency across three climatically diverse testbeds

Metric	F-statistic	p-value	$\eta^2$ (eta-squared)	Effect Size	Regional Differences
Solar forecast error (nRMSE)	329.45	<0.001	0.0245	Medium	Yes
Wind forecast error (nRMSE)	51.86	<0.001	0.0039	Small	Yes
Voltage compliance (%)	1.11	0.329	0.0001	Small	No
Renewable penetration (%)	—	<0.001	1.0000	Large	Yes
Solar agent response time (ms)	0.55	0.580	0.0000	Small	No
Wind agent response time (ms)	51.86	<0.001	0.0039	Small	Yes

**Notes:**

- One-way ANOVA conducted across three regions (Morocco, Denmark, UAE) with  $n=8,760$  per region
- $\eta^2$  (eta-squared) represents proportion of total variance explained by regional differences
- Effect size interpretation: Small ( $\eta^2 < 0.06$ ), Medium ( $0.06 \leq \eta^2 < 0.14$ ), Large ( $\eta^2 \geq 0.14$ )
- Significance threshold:  $p < 0.05$  indicates statistically significant regional differences
- Solar and wind forecast errors measured via Equations 2 and 6 respectively
- Voltage compliance evaluated via Equation 19, renewable penetration via Equation 22, response time via Equation 16
- "Yes" in Regional Differences column indicates significant performance heterogeneity requiring climate-specific tuning
- "No" indicates framework robustness across all three climatically diverse testbeds

Table 9 provides analysis of variance across three climatically diverse regional testbeds, quantifying the extent to which geographic and climatic factors influence MARL-MADDPG framework performance. Solar forecasting demonstrates statistically significant regional differences ( $F=329.45$ ,  $p < 0.001$ ) with medium effect size ( $\eta^2=0.0245$ ), indicating that 2.45% of forecasting error variance stems from climatic heterogeneity rather than stochastic renewable variability, as calculated via Equation 2. This medium effect validates the necessity of climate-specific ARIMA model tuning for optimal solar prediction across desert, maritime, and extreme thermal environments.

Wind forecasting exhibits significant yet smaller regional variation ( $F=51.86$ ,  $p < 0.001$ ,  $\eta^2=0.0039$ ), with only 0.39% of error variance attributable to geographic factors, confirming LSTM neural network robustness to atmospheric turbulence patterns across diverse climates when evaluated via Equation 6. Voltage compliance demonstrates no significant regional differences ( $F=1.11$ ,  $p=0.329$ ), validating distributed consensus protocol effectiveness (Equations 17 and 18) regardless of grid infrastructure characteristics or renewable penetration levels.

Renewable penetration reveals the most substantial regional heterogeneity ( $\eta^2=1.0000$ , large effect size,  $p < 0.001$ ) driven by Denmark's 294.7% offshore wind capacity versus Morocco and UAE's solar-dominated profiles (148.8% and 134.9% respectively), as quantified via Equation 22. This large effect reflects fundamental resource availability differences rather than framework limitations.

Agent response times demonstrate climate-independent performance for solar agents ( $F=0.55$ ,  $p=0.580$ ) but statistically significant yet practically negligible variations for wind agents ( $F=51.86$ ,  $p < 0.001$ ,  $\eta^2=0.0039$ ), with all response times remaining within 42 to 62 millisecond ranges satisfying IEEE 1547-2018 requirements via Equation 16. The small effect sizes confirm framework transferability across climatically diverse deployment contexts.

#### 4.3.2 Scalability analysis and resilience validation under extreme conditions

Computational scalability constitutes a critical determinant of practical deployment viability, as distributed MARL frameworks must maintain real time performance across varying agent populations while managing resource constraints inherent in edge computing infrastructure. Scalability evaluation quantifies computational resource utilization CPU, memory, response latency, network bandwidth as agent populations increases from minimal 4 agent

configurations to large scale 100 agent deployments representative of utility scale microgrids. Complementary resilience testing validates framework robustness under operational disruptions including agent failures, communication degradation, and extreme forecasting errors that characterize real world deployment environments. Table 10 presents computational resource scaling characteristics across six agent population configurations, establishing practical deployment boundaries for edge computing platforms.

**Table 10.** Computational resource scaling and performance degradation across varying agent populations

Agent Population	CPU Utilization (%)	Memory Footprint (GB)	Response Time (ms)	Network Bandwidth (Mbps)	Coordination Efficiency (%)
4	38.2 ± 2.1	0.9	41.3 ± 3.2	1.2	96.8
10	46.5 ± 2.4	2.4	45.8 ± 3.8	2.4	95.2
25	58.1 ± 3.1	4.7	52.1 ± 4.5	8.7	92.7
50	74.3 ± 3.8	9.2	67.4 ± 6.2	23.5	88.4
75	82.1 ± 4.2	14.3	82.7 ± 8.1	52.8	84.1
100	89.7 ± 5.1	18.6	94.2 ± 9.8	87.3	79.3

**Notes:**

- CPU Utilization and Memory Footprint measured on Apple MacBook M1 Max hardware (32 core GPU, 32 GB unified memory) during continuous 8,760 hour operational simulation
- Response Time quantifies mean agent decision latency via Equation 16 across all agent types (solar, wind, storage, load) with 95% confidence intervals
- Network Bandwidth represents inter agent communication overhead for distributed consensus protocols (Equations 17 to 18) under fully connected topology
- Coordination Efficiency measures percentage of optimal joint reward  $R_{\text{joint}}$  achieved relative to centralized omniscient baseline
- Agent populations 4 to 25 represent distribution level microgrids, 50 to 75 utility scale installations, 100 large interconnected systems
- Critical threshold approximately 75 agents where response time approaches 100 millisecond IEEE 1547 2018 real time limit
- Computational complexity scales as  $O(N \cdot \log K)$  for decentralized execution where  $N$  is agent count and  $K$  is communication neighbors
- All metrics averaged across Morocco, Denmark, and UAE testbeds with standard deviations representing cross regional variability

Table 10 demonstrates that the distributed MADDPG framework maintains computational tractability across agent populations spanning two orders of magnitude. CPU utilization exhibits sublinear growth from 38.2% at 4 agents to 89.7% at 100 agents, confirming logarithmic complexity scaling predicted by distributed algorithm theory rather than polynomial growth characteristic of centralized optimization. Memory footprint increases approximately linearly from 0.9 GB to 18.6 GB, reflecting experience replay buffer storage requirements proportional to agent count, suggesting practical deployment limits around 75 to 100 agents on commercial edge devices with 16 to 32 GB RAM capacity. Response time remains below the critical 100 millisecond threshold up to 75 agents (82.7 milliseconds), beyond which latency constraints necessitate distributed computing infrastructure with shared replay buffers or sparse communication topologies reducing neighbor connectivity. Network bandwidth scales quadratically in fully connected topologies from 1.2 Mbps at 4 agents to 87.3 Mbps at 100 agents, indicating that microgrids exceeding 50 agents require hierarchical communication architectures limiting each agent to 5 to 8 neighbors rather than global connectivity. Coordination efficiency degradation from 96.8% to 79.3% reflects increased consensus protocol convergence time as agent populations grow, though performance remains substantially above independent learning baselines (Table 6) even at maximum scale.

Beyond scalability constraints, operational resilience under component failures and extreme conditions determines framework suitability for critical infrastructure deployment where unplanned outages, communication disruptions, and meteorological anomalies represent routine operational challenges. Resilience testing quantifies performance degradation, recovery characteristics, and autonomous adaptation capabilities when subjected to single agent failures, communication latency spikes, and extreme forecasting errors exceeding 97th percentile thresholds. Table 11 presents comprehensive failure scenario analysis across five operational disruption categories validated through 26,280 simulated operational hours.

**Table 11.** Performance degradation and recovery under operational failure scenarios and extreme conditions

Failure Scenario	Performance Degradation (%)	Recovery Time (s)	Voltage Compliance (%)	Response Time (ms)	System Uptime (%)
Single solar agent dropout	10.2 ± 1.8	3.2 ± 0.8	95.3	47 ± 5	100.0
Single wind agent dropout	11.7 ± 2.1	3.5 ± 0.9	94.8	48 ± 6	100.0
Single storage agent dropout	16.4 ± 2.7	4.1 ± 1.2	92.7	52 ± 8	100.0
Communication latency 200 ms (5% duration)	7.3 ± 1.4	N/A	97.0	78 ± 12	100.0
Extreme forecast error (nRMSE >8%)	12.8 ± 2.3	Adaptive	94.2	51 ± 7	100.0
Baseline (no failures)	0.0	N/A	97.2	45 ± 4	100.0

**Notes:**

- Performance Degradation measured as percentage reduction in joint reward  $R_{\text{joint}}$  (Equation 26) relative to baseline operation without failures
- Recovery Time quantifies duration for distributed consensus protocols (Equations 17 to 18) to redistribute agent responsibilities following sudden dropout
- Voltage Compliance percentage of operational hours maintaining IEEE 1547 2018 limits (0.95 to 1.05 per unit) via Equation 19
- Response Time measured via Equation 16 reflecting increased computational load during failure compensation
- System Uptime percentage of 8,760 hour validation period maintaining operational functionality without complete system failure
- Single agent dropout scenarios simulate unannounced permanent failure without graceful shutdown or handover protocols
- Communication latency 200 millisecond delays imposed for 5% of operational periods (438 hours annually) simulating network congestion
- Extreme forecast error scenarios represent 97th percentile events (nRMSE >8%) occurring during sandstorms (Morocco, UAE) and rapid atmospheric fronts (Denmark), totaling approximately 263 hours annually per testbed
- Adaptive recovery for forecasting errors indicates continuous autonomous adjustment without discrete recovery event
- All metrics averaged across Morocco, Denmark, and UAE testbeds with ± representing standard deviations across regional variations

Table 11 validates robust fault tolerance characteristics essential for autonomous critical infrastructure deployment. Single agent dropout scenarios demonstrate graceful degradation rather than cascading failures, with performance reductions ranging from 10.2% for solar agents to 16.4% for storage agents, the latter reflecting critical arbitrage role

in renewable surplus absorption and demand peak coverage. Rapid recovery times of 3.2 to 4.1 seconds confirm distributed consensus protocol effectiveness (Equations 17 to 18) in redistributing load balancing responsibilities through neighboring agent coordination without centralized intervention. Communication latency spikes imposing 200 millisecond delays for 5% of operational periods increase mean response times from baseline 45 milliseconds to 78 milliseconds, remaining below the 100 millisecond IEEE 1547 2018 real time threshold while maintaining 100% system uptime and 97.0% voltage compliance. Extreme forecasting error scenarios exceeding 97th percentile thresholds trigger adaptive storage reserve margin increases from baseline 0.77 average state of charge to 0.85, maintaining voltage compliance at 94.2% compared to 97.2% baseline through autonomous resilience mechanisms. Critically, no failure scenario precipitated complete system collapse, validating the distributed architecture advantage over centralized single point of failure vulnerabilities. The consistent 100% system uptime across all disruption categories, contrasted against centralized MPC vulnerabilities documented in prior microgrid reliability studies, substantiates the distributed MARL framework suitability for mission critical renewable energy infrastructure requiring fault tolerant autonomous operation.

## 5. Discussion

This research establishes a simulation based artificial intelligence framework employing multi agent reinforcement learning for autonomous hybrid microgrid control, validated across 26,280 operational hours spanning three climatically diverse regions representing Morocco desert solar conditions with sandstorm variability, Denmark maritime offshore wind environments, and UAE extreme thermal stress with dust related intermittency. By integrating MADDPG actor critic neural networks with hidden layer architecture alongside specialized ARIMA solar forecasting and LSTM wind prediction agents, the framework demonstrates distributed coordination under renewable uncertainty, providing superior performance relative to centralized model predictive control and independent single agent reinforcement learning approaches. The validated forecasting accuracy demonstrates that integrated LSTM-ARIMA-MADDPG coordination enables distributed agents to adapt prediction models to regional climatic conditions while maintaining real-time computational feasibility. This performance validates the framework design principle that specialized neural network architectures for heterogeneous renewable resources improve coordination effectiveness compared to generic forecasting baselines, corroborating findings by Zhu et al. and Hadi et al. [72,73] while extending validation across substantially longer horizons and more diverse geographic contexts.

Beyond aggregate forecasting performance, multi-regional validation reveals critical insights into framework transferability and resource driven performance boundaries. While neural network execution latency response times 45.3 to 45.6 milliseconds, coordination efficiency remains invariant across climates with negligible effect sizes  $d < 0.1$  despite 8,760 hour validation periods per region, validating that MADDPG policy optimization generalizes independent of geographic context, renewable penetration outcomes exhibit substantial regional heterogeneity. Denmark  $2.1\times$  advantage over Morocco in renewable integration 291.7% versus 140.6% stems fundamentally from offshore wind resource availability rather than algorithmic superiority, establishing that distributed MARL coordination optimizes available resources but cannot overcome physical generation constraints. This distinction proves critical for deployment planning where operators in resource limited regions should prioritize capacity expansion alongside MARL implementation rather than expecting algorithmic optimization alone to achieve high penetration targets. Variance decomposition analysis clarifies that only 9.5% of renewable penetration variance stems from regional differences while 90.5% reflects within region operational dynamics, whereas energy efficiency 32.3% regional variance component indicates thermal conditions impose hard physical constraints on turbine and photovoltaic performance that MARL coordination cannot eliminate, only optimize within. These patterns align with Ramesh et al. findings demonstrating advanced reinforcement learning methods outperforming traditional approaches in renewable energy management [74], while Li et al. extended understanding through federated MARL applications highlighting superior coordination and data privacy in distributed energy systems [75].

System uptime reached 100.0% across all testbeds with voltage maintained within IEEE 0.95 to 1.05 per unit limits, eliminating single point of failure vulnerabilities characteristic of centralized topologies, though this 100% voltage

compliance achieved across 26,280 simulated operational hours will likely degrade to 92 to 96% compliance in hardware implementations due to sensor measurement noise  $\pm 2\%$  typical for voltage transducers, actuator response delays 50 to 150 millisecond mechanical switching, and communication jitter inherent in wireless 5G networks, as documented in operational microgrid studies demonstrating 3 to 5% performance gaps between simulation and physical deployment. Load balancing efficiency exhibited regional variation Morocco 82.1%, Denmark 85.6%, UAE 87.8%, slightly below 90% targets but reflecting realistic constraints under variable renewable generation, with communication success rates of 99.8 to 99.9% confirming robust distributed coordination validated by Khan and Wang [76] comprehensive survey on multi agent systems for microgrid control and Nakabi and Toivanen reporting substantial energy management effectiveness gains through deep reinforcement learning while preserving system stability [46]. To contextualize simulation performance against operational systems, comparison with established microgrids including Stone Edge Farm California 160 kW photovoltaic with hydrogen fuel cells achieving near 100% renewable operation with millisecond scale control [77], Bordsesholm Germany demonstrating 100% renewable supply from integrated photovoltaic biomass combined heat power [78], and Kythnos Greece operational since 2001 with 10 kWp photovoltaic and 53 kWh battery storage suggests the simulated MARL framework demonstrates performance comparable to or exceeding these operational benchmarks [79]. However, this comparison requires critical qualification as real world microgrids encounter hardware constraints sensor drift, actuator delays, communication packet loss, regulatory compliance requirements grid code adherence, interconnection protocols, and operational uncertainties extreme weather events, equipment degradation, cybersecurity threats abstracted in the present simulation framework, with hardware in the loop validation and pilot deployments likely revealing 10 to 25% performance degradation relative to simulation results due to these cyber physical realities [19,46].

Scalability analysis establishes practical deployment boundaries absent from prior MARL energy research through systematic evaluation across agent populations spanning 4 to 100 units Table 10. CPU utilization exhibits sublinear growth from 38.2% at 4 agents to 89.7% at 100 agents, confirming logarithmic complexity scaling  $O(N \cdot \log K)$  predicted by distributed algorithm theory rather than polynomial growth characteristic of centralized optimization, while memory footprint increases approximately linearly from 0.9 GB to 18.6 GB reflecting experience replay buffer storage requirements proportional to agent count. Commercial edge devices with 8 to 16 GB RAM support approximately 75 to 100 agents before requiring costly distributed computing infrastructure with shared replay buffers across multiple nodes, establishing hard infrastructure constraints particularly relevant for utility scale microgrids spanning multiple feeders or substations where agent populations could exceed 100 distributed across 50 generation units, 30 storage systems, and 40 load zones. Response time remains below critical 100 millisecond IEEE 1547 2018 threshold up to 75 agents 82.7 milliseconds, beyond which latency constraints necessitate distributed computing infrastructure or sparse communication topologies reducing neighbor connectivity. Network bandwidth scales quadratically in fully connected topologies from 1.2 Mbps at 4 agents to 87.3 Mbps at 100 agents, indicating microgrids exceeding 50 agents require hierarchical communication architectures limiting each agent to 5 to 8 neighbors rather than global connectivity. Complementary resilience testing Table 11 validates robust fault tolerance characteristics essential for autonomous critical infrastructure deployment, with single agent dropout scenarios demonstrating graceful degradation rather than cascading failures where performance reductions range from 10.2% for solar agents to 16.4% for storage agents, the latter reflecting critical arbitrage role in renewable surplus absorption and demand peak coverage. Rapid recovery times of 3.2 to 4.1 seconds confirm distributed consensus protocol effectiveness Equations 17 to 18 in redistributing load balancing responsibilities through neighboring agent coordination without centralized intervention. However, coordinated cyberattacks targeting multiple agents simultaneously 3+ compromised in 10 agent systems could cause cascading voltage control failures, necessitating Byzantine resilient consensus protocols employing cryptographic verification and outlier detection [46]. Operational 5G networks experience 5 to 10% packet loss under congestion versus simulated  $<2\%$  assumption, requiring redundant communication paths dual 5G fiber, message acknowledgment protocols with exponential backoff retransmission, and local autonomy modes enabling 30 to 60 second operation without coordination during outages [22,63].

When compared to centralized model predictive control baselines ranging from 57.6 to 74.0 dollars per MWh, the multi agent framework achieved LCOE values of 50.7 to 65.1 dollars per MWh, yielding 12% cost reduction more conservative than 20 to 30% savings reported in prior research due to accounting for computational overhead, communication latencies, and stochastic renewable variability often omitted in idealized studies. However, the 12% LCOE advantage assumes centralized MPC operational expenditure remains constant over decade scale horizons, whereas centralized systems may incorporate advanced forecasting improvements and machine learning enhanced optimization potentially narrowing MARL advantage to 5 to 8% in commercial deployments. Cumulative ten year economic benefits projected \$900 million Morocco, \$1,000 million Denmark, and \$1,100 million UAE under stated modeling assumptions, though realistic projections likely exhibit non-linear trajectories with steeper benefits in years 1 to 5 followed by diminishing returns due to battery capacity fade reducing arbitrage revenue, technology learning curves benefiting both architectures equally, and regulatory changes carbon pricing altering relative economics. The 5% discount rate may underestimate risk premiums in developing markets Morocco, UAE where financing costs for AI based control systems could reach 8 to 12%, reducing net present value by 25 to 40%. Technology adoption forecasts suggesting 85 to 95% MARL penetration by 2030 reflect optimistic scenarios, as actual adoption depends on regulatory approval requiring 12 to 18 month utility validation periods before commercial operation authorization, workforce training programs for SCADA operator transitions to AI driven control, insurance industry acceptance of autonomous systems liability frameworks not yet established in most jurisdictions, and multi-year pilot validation demonstrating reliability under extreme conditions.

Training dynamics validated MADDPG framework reproducibility with convergence thresholds reached at episodes 6,200 Denmark, 6,500 Morocco, and 9,800 UAE reflecting region specific learning complexities, demonstrating 9 to 15% performance superiority over Independent Q Learning baselines confirming centralized critic advantages for multi agent credit assignment. Hyperparameter sensitivity analysis identified optimal configurations learning rate  $\alpha = 0.001$ , batch size  $B = 64$ , discount factor  $\gamma = 0.95$  achieving 88.0% performance stability with narrow standard deviations  $\pm 0.5$  to 1.0% Table 7, establishing reproducible parameter selection protocols validated across 26,280 operational hours spanning three regional testbeds. MADDPG training converged within 18 to 22 hours on Apple MacBook M1 Max hardware 32 core GPU, 32 GB unified memory Figure 6 Panel B, with CPU utilization between 46.0 and 46.9% and memory usage from 62.8 to 63.2% demonstrating computational efficiency compatible with commercial edge computing platforms. These moderate resource demands remove specialized high performance infrastructure requirements and lower capital deployment barriers for validated 4 to 75 agent populations, though larger utility scale deployments exceeding 75 agents require distributed infrastructure investments with shared replay buffers increasing capital costs by 40 to 60%. Future deployment on industrial edge computing platforms may necessitate neural network model compression techniques to meet cost and power constraints while maintaining sub 100 millisecond response time requirements for real time grid control.

These computational characteristics translate into several practical implications for utilities, regulators, and researchers that guide how the framework can be moved from simulation to deployment.

- Utility operators and microgrid developers: Standardized MARL training protocols enable workforce transition from conventional SCADA based operation to AI driven autonomous control. Commercial edge devices with 16 to 32 GB RAM can host 75 to 100 agents, which is suitable for distribution level microgrids spanning multiple feeders, while one time training times of 18 to 22 hours on commodity GPU hardware remain compatible with commissioning schedules.
- Communication and infrastructure planning: Network latency around 31.3 milliseconds and communication success rates exceeding 99.8 percent indicate feasibility over industrial 5G networks, but electromagnetic interference in practical environments is likely to increase packet loss from the simulated 2 percent toward 5 to 10 percent. This motivates redundant communication architectures, such as dual 5G and fiber links with acknowledgment and retransmission mechanisms, before field deployment.
- Geographical deployment and regulation: Demonstrated robustness across desert Morocco, maritime Denmark, and extreme heat UAE suggests applicability to a wide range of renewable portfolios, from North

African solar to Scandinavian offshore wind. However, large scale deployment still requires hardware in the loop validation to account for sensor drift, actuator delays, and country specific grid codes and interconnection procedures.

- Economic decision making: The simulated 12 percent LCOE reduction and ten year cumulative benefits on the order of hundreds of millions of dollars per region provide an initial business case for utilities, but conservative margins are necessary when planning investments because simulation to field performance gaps and financing conditions in different markets can reduce realized gains.

## 6. Conclusion

This research developed and validated an artificial intelligence framework employing multi-agent reinforcement learning to achieve autonomous control in hybrid renewable energy microgrids through simulation-based analysis. The framework integrated three neural network architectures for distributed coordination. Long Short-Term Memory (LSTM) recurrent neural networks with memory cell structures for wind power forecasting, Multi-Agent Deep Deterministic Policy Gradient (MADDPG) actor-critic deep neural networks with progressive dimensionality reduction through \*\*\*\* hidden layer architecture for battery storage dispatch optimization, and attention-based centralized critic networks enabling dynamic neighboring agent coordination. This system underwent validation across 26,280 operational hours within the geographically distinct contexts of Morocco, Denmark, and the United Arab Emirates. The simulation demonstrated that distributed agent architectures could maintain grid stability, achieve cost-competitive operation, and adapt to heterogeneous climatic conditions ranging from desert solar environments to maritime wind-dominated systems and extreme heat scenarios.

The principal contribution lies in establishing a complete AI-driven control paradigm that transcends conventional centralized optimization through validated distributed intelligence mechanisms. Unlike prior MARL applications limited to theoretical benchmarks or short-duration tests, this framework delivers 26,280-hour operational validation (8,760 hours per region) with quantified performance across technical, economic, and reliability metrics simultaneously. By combining LSTM recurrent neural networks for wind prediction processing sequential atmospheric data, MADDPG actor-critic deep neural networks with \*\*\*\* architecture for storage arbitrage, and ARIMA for solar forecasting within a unified multi-agent architecture, the framework addresses the fundamental heterogeneity challenge in renewable resource prediction through specialized neural network implementations. This methodological advancement results in significant error reductions of 28.3% for solar and 31.8% for wind compared to single method baselines, while storage prediction accuracy increases by 37.5%. Systematic hyperparameter optimization through 80-trial grid search established validated neural network training configurations (learning rate 0.001, batch size 64, discount factor 0.95) achieving 88.0% performance stability with narrow standard deviations ( $\pm 0.5$  to 1.0%), providing reproducible implementation protocols absent in existing literature. Geographic transferability is demonstrated in simulation through the framework's performance across diverse renewable portfolios in Morocco, Denmark, and the UAE. Despite the variations between solar dominated, wind intensive, and desert solar configurations, the system maintained 100% voltage compliance and outperformed Independent Q Learning baselines by 9 to 15%. Such results substantiate the potential for global commercial deployment, moving beyond the limitations of earlier single region studies.

As a simulation-based investigation, several operational complexities remain outside the current scope. The main limitations are:

- Communication modeling: Idealized communication is assumed without realistic packet loss, latency spikes, or cyber intrusions, whereas operational networks can experience higher loss rates and security constraints.
- Hardware abstraction: Physical effects such as thermal degradation of edge processors, sensor drift in measurement channels, and actuator delays in switching devices are not explicitly modeled, even though they can degrade control performance in practice.

- Regulatory and operational context: Utility interconnection procedures, grid code certification, cybersecurity compliance, and insurance or liability frameworks for autonomous controllers are not represented in the simulations.
- Economic assumptions: Projected cost reductions depend on the persistence of energy price differentials and assume that centralized model predictive control does not improve at the same rate, which may not hold over multi-year horizons.

These limitations define appropriate methodological boundaries for algorithm development and multi-regional validation. Hardware in the loop experimentation and pilot deployments represent essential next steps before large scale commercial implementation.

Future investigations should prioritize hardware-in-the-loop validation integrating physical microgrid components with simulated agents to assess real-time control performance under actual hardware constraints. Expanding the study to additional climate zones, such as tropical monsoon regions, Arctic environments, and high-altitude installations, would further strengthen generalizability claims. Furthermore, Integration with edge AI accelerators and neuromorphic computing platforms could enhance computational efficiency for large-scale deployments, while federated learning protocols would enable privacy-preserving multi-stakeholder coordination across interconnected microgrids. Finally, resilience testing against cyberattacks and extreme weather remains essential for validating the framework in critical infrastructure applications.

Overall, this validated AI framework demonstrates potential to advance microgrid control from reactive energy balancing to predictive autonomous operation. Multi agent reinforcement learning employing LSTM recurrent neural networks and MADDPG actor critic deep neural network architectures emerge as a promising simulation validated approach that now requires hardware validation for intelligent renewable energy systems. The demonstrated feasibility of distributed artificial intelligence managing complex power networks under realistic operational constraints directly supports the renewable energy transition, offering a scalable pathway from simulation environments to commercial deployment in pursuit of decarbonized, autonomously coordinated electrical grids.

## Appendix A. Nomenclature and Abbreviations

**Table A1.** Complete nomenclature and abbreviations used in this study

Abbreviation	Full Form
<b>5G</b>	Fifth Generation Wireless Communication
<b>AC</b>	Alternating Current
<b>AI</b>	Artificial Intelligence
<b>ANOVA</b>	Analysis of Variance
<b>API</b>	Application Programming Interface
<b>ARIMA</b>	AutoRegressive Integrated Moving Average
<b>BCa</b>	Bias-Corrected and Accelerated
<b>CAPEX</b>	Capital Expenditure
<b>CPU</b>	Central Processing Unit
<b>CSP</b>	Concentrated Solar Power
<b>CTDE</b>	Centralized Training with Decentralized Execution
<b>DC</b>	Direct Current
<b>DDPG</b>	Deep Deterministic Policy Gradient
<b>GB</b>	Gigabyte
<b>GPU</b>	Graphics Processing Unit
<b>HVAC</b>	Heating, Ventilation, and Air Conditioning
<b>IEEE</b>	Institute of Electrical and Electronics Engineers
<b>IoT</b>	Internet of Things

<b>IQL</b>	Independent Q-Learning
<b>kW</b>	Kilowatt
<b>kWh</b>	Kilowatt-hour
<b>kWh/m<sup>2</sup>/year</b>	Kilowatt-hour per square meter per year
<b>LCOE</b>	Levelized Cost of Electricity
<b>LSTM</b>	Long Short-Term Memory
<b>MADDPG</b>	Multi-Agent Deep Deterministic Policy Gradient
<b>MARL</b>	Multi-Agent Reinforcement Learning
<b>MPC</b>	Model Predictive Control
<b>ms</b>	Millisecond
<b>MW</b>	Megawatt
<b>MWh</b>	Megawatt-hour
<b>NASA</b>	National Aeronautics and Space Administration
<b>nRMSE</b>	Normalized Root Mean Square Error
<b>OPEX</b>	Operational Expenditure
<b>p.u.</b>	Per Unit
<b>POWER</b>	Prediction Of Worldwide Energy Resources
<b>PV</b>	Photovoltaic
<b>ReLU</b>	Rectified Linear Unit
<b>RL</b>	Reinforcement Learning
<b>ROI</b>	Return on Investment
<b>SCADA</b>	Supervisory Control and Data Acquisition
<b>SD</b>	Standard Deviation
<b>SOC</b>	State of Charge
<b>UAE</b>	United Arab Emirates
<b>USD</b>	United States Dollar

**Acknowledgments:** This research was supported by the “University of Debrecen Program for Scientific Publication”.

**Conflicts of Interest:** The author declares no conflicts of interest.

**Author Contributions:** Conceptualization, M.F.R.; methodology, M.F.R.; software, M.F.R.; validation, M.F.R.; formal analysis, M.F.R.; investigation, M.F.R.; resources, M.F.R.; data curation, M.F.R.; writing—original draft preparation, M.F.R.; writing—review and editing, M.F.R.; visualization, M.F.R.; project administration, M.F.R. The author has read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable. This study is based on mathematical simulations and computational modeling and does not involve human subjects or animal experimentation.

**Informed Consent Statement:** Not applicable. This study did not involve human participants.

#### **Data Availability Statement**

All findings are based on computational simulation of multi agent reinforcement learning (MADDPG) for hybrid microgrid control across Morocco, Denmark, and the UAE. The complete dataset comprises fifteen CSV files containing hourly operational metrics, AI training convergence, statistical validation, and performance benchmarks, totaling 26,280 simulated operational hours across three regional testbeds.

## Simulation Outputs

### AI training and sensitivity analysis:

1. marl\_training\_convergence.csv – Actor/critic loss, joint reward, and Q value evolution over 10,000 training episodes for solar, wind, and storage agents
2. marl\_hyperparameter\_sensitivity.csv – Performance under varied learning rates, batch sizes, and discount factors
3. marl\_algorithm\_comparison.csv – Comparative analysis of MADDPG, PPO, SAC, and Independent Q Learning
4. marl\_neural\_network\_analysis.csv – Layer wise activations, gradient flow, and weight update statistics

### Statistical validation and reproducibility analysis:

5. marl\_training\_multiple\_runs\_raw.csv – Multiple independent training runs with different random seeds for reproducibility validation
6. marl\_training\_statistics\_with\_ci.csv – Training statistics with 95% confidence intervals across multiple runs
7. marl\_failure\_scenarios.csv – Agent failure resilience testing and recovery performance metrics
8. marl\_scalability\_analysis.csv – Computational resource utilization across varying agent populations (4 to 100 agents)
9. marl\_training\_stability.csv – Convergence stability metrics and variance analysis across regional testbeds
10. marl\_statistical\_significance\_tests.csv – Pairwise comparison statistics and effect size analysis

### Regional operational data:

11. marl\_operational\_morocco.csv – Hourly PV/wind generation, storage SOC, voltage stability, forecasting errors, and agent response times (8,760 hours)
12. marl\_operational\_denmark.csv – Hourly operational metrics for Denmark (8,760 hours)
13. marl\_operational\_uae.csv – Hourly operational metrics for the UAE (8,760 hours)
14. marl\_operational\_multi\_regional.csv – Consolidated dataset combining all three regions (26,280 hours)
15. marl\_performance\_summary.csv – Aggregated forecasting accuracy (nRMSE), LCOE, renewable penetration, and grid stability metrics by region

## Data Sources and Reproducibility

Regional renewable resource profiles and meteorological boundary conditions were derived from publicly accessible datasets:

- NASA Prediction of Worldwide Energy Resources (POWER). National Aeronautics and Space Administration Langley Research Center. Available at: <https://power.larc.nasa.gov> (accessed August 11, 2025).
- Global Wind Atlas. DTU Wind Energy and World Bank Group. Available at: <https://globalwindatlas.info> (accessed August 11, 2025).
- National Electricity Regulatory Authority (ANRE), Morocco. Available from Ministry of Energy Transition and Sustainable Development.
- Danish Energy Agency. Energy Statistics 2023. Available at: <https://ens.dk/en/analyses-and-statistics/key-figures> (accessed August 11, 2025).
- Ministry of Energy and Infrastructure, UAE. Electricity & Renewable Energy Open Data. Available at: <https://opendata.moei.gov.ae/topic/electricity-renewable-energy> (accessed August 11, 2025).

**Declaration of Generative AI and AI-assisted Technologies:** During preparation of this manuscript, DeepL Write and ScienceDirect AI were utilized exclusively for language refinement, grammar correction, academic tone enhancement, and manuscript formatting. DeepL Write provided sentence-level writing improvements and clarity enhancements. ScienceDirect AI assisted with literature search refinement and related terminology accuracy. All research design, methodology, data analysis, result interpretation, and scientific conclusions were entirely conceived, executed, and verified by the author. Following application of these tools, the manuscript was thoroughly reviewed and edited to ensure all content reflects the author's original work and scientific intent. The author assumes full responsibility for the accuracy, integrity, and scientific validity of this publication.

## References

- [1] R. Xue, G. Zhang, Z. Li, Home Energy Systems in Europe: Advancements and Future Directions, *Chinese Journal of Electrical Engineering* 11 (2025) 38–62. <https://doi.org/10.23919/CJEE.2025.000137>.
- [2] S. Tan, P. Xie, J.C. Vasquez, J.M. Guerrero, Developments, challenges and future opportunities in cybersecure microgrid control, *Nature Reviews Electrical Engineering* (2025). <https://doi.org/10.1038/s44287-025-00189-z>.
- [3] A. Parizad, H.R. Baghaee, V. Alizadeh, S. Rahman, Emerging Technologies and Future Trends in Cyber-Physical Power Systems: Toward a New Era of Innovations, in: *Smart Cyber-Physical Power Systems*, Wiley, 2025: pp. 525–565. <https://doi.org/10.1002/9781394334599.ch19>.
- [4] M.F. Rabbi, A machine learning framework for forecasting multidimensional sustainability and informing integrated policy thresholds in the EU, *Environ Dev Sustain* (2025). <https://doi.org/10.1007/s10668-025-06762-8>.
- [5] H. Javed, F. Eid, S. El-Sappagh, T. Abuhmed, Sustainable energy management in the AI era: a comprehensive analysis of ML and DL approaches, *Computing* 107 (2025) 132. <https://doi.org/10.1007/s00607-025-01485-0>.
- [6] S. Dong, C. Wang, Y. Zhang, Y. Wang, X. Zhang, L. Guo, L. Ju, Hierarchical deep Q-network-based optimization of resilient grids under multi-dimensional uncertainties from extreme weather, *Sci Rep* 15 (2025) 24927. <https://doi.org/10.1038/s41598-025-09868-1>.
- [7] K. Ronay, D. Bica, C. Munteanu, Micro-grid Development Using Artificial Neural Network for Renewable Energy Forecast and System Control, *Procedia Eng* 181 (2017) 818–823. <https://doi.org/10.1016/j.proeng.2017.02.472>.
- [8] M.F. Rabbi, Unified artificial intelligence framework for modeling pollution dynamics and sustainable remediation in environmental chemistry, *Sci Rep* 15 (2025) 36196. <https://doi.org/10.1038/s41598-025-20083-w>.
- [9] R. Benbba, M. Barhdadi, A. Ficarella, G. Manente, M.P. Romano, N. El Hachemi, A. Barhdadi, A. Al-Salaymeh, A. Outzourhit, Solar Energy Resource and Power Generation in Morocco: Current Situation, Potential, and Future Perspective, *Resources* 13 (2024) 140. <https://doi.org/10.3390/resources13100140>.
- [10] P. Fragkos, Assessing the energy system impacts of Morocco's nationally determined contribution and low-emission pathways, *Energy Strategy Reviews* 47 (2023) 101081. <https://doi.org/10.1016/j.esr.2023.101081>.
- [11] A.Q. Santos, Z. Ma, C.G. Olsen, B.N. Jørgensen, Framework for Microgrid Design Using Social, Economic, and Technical Analysis, *Energies (Basel)* 11 (2018) 2832. <https://doi.org/10.3390/en11102832>.

- [12] R. McKenna, M. D'Andrea, M.G. González, Analysing long-term opportunities for offshore energy system integration in the Danish North Sea, *Advances in Applied Energy* 4 (2021) 100067. <https://doi.org/10.1016/j.adapen.2021.100067>.
- [13] A. AlHammadi, N. Al-Saif, A.S. Al-Sumaiti, M. Marzband, T. Alsumaiti, E. Heydarian-Forushani, Techno-Economic Analysis of Hybrid Renewable Energy Systems Designed for Electric Vehicle Charging: A Case Study from the United Arab Emirates, *Energies (Basel)* 15 (2022) 6621. <https://doi.org/10.3390/en15186621>.
- [14] A. Alraeesi, A.H. Shah, A. Hassan, M.S. Laghari, Characterisation of Dust Particles Deposited on Photovoltaic Panels in the United Arab Emirates, *Applied Sciences* 13 (2023) 13162. <https://doi.org/10.3390/app132413162>.
- [15] M.F. Rabbi, Sustainable Reduced-Order Thermal Modeling for Energy-Efficient Real-Time Control of Grid-Scale Energy Storage Systems, *Sustainability* 17 (2025) 9839. <https://doi.org/10.3390/su17219839>.
- [16] I. Ahmed, M.A. Syed, M. Maaruf, M. Khalid, Distributed computing in multi-agent systems: a survey of decentralized machine learning approaches, *Computing* 107 (2025) 2. <https://doi.org/10.1007/s00607-024-01356-0>.
- [17] Q.L. Lam, D. Riu, A.I. Bratcu, Frequency Robust Control Application in Islanded Microgrids Considering Parametric Uncertainties and Distinct Photovoltaic Penetration Rate Scenarios, *IEEE Access* 11 (2023) 92589–92616. <https://doi.org/10.1109/ACCESS.2023.3307210>.
- [18] Q.L. Lam, D. Riu, A.I. Bratcu, A. Labonne, C. Boudinet, Power hardware-in-the-loop validation of primary frequency robust control in stand-alone microgrids with storage units, *Electrical Engineering* 105 (2023) 317–333. <https://doi.org/10.1007/s00202-022-01666-6>.
- [19] Q.L. Lam, A.I. Bratcu, D. Riu, Multi-Variable  $H_{\infty}$  Control Approach for Voltage Ancillary Service in Autonomous Microgrids: Modelling, Design, and Sensitivity Analysis, *IEEE Access* 9 (2021) 140212–140234. <https://doi.org/10.1109/ACCESS.2021.3119375>.
- [20] Q.L. Lam, A.I. Bratcu, D. Riu, C. Boudinet, A. Labonne, M. Thomas, Primary frequency control in stand-alone microgrids with storage units: A robustness analysis confirmed by real-time experiments, *International Journal of Electrical Power & Energy Systems* 115 (2020) 105507. <https://doi.org/10.1016/j.ijepes.2019.105507>.
- [21] A.M. Bouzid, J.M. Guerrero, A. Cheriti, M. Bouhamida, P. Sicard, M. Benhanem, A survey on control of electric power distributed generation systems for microgrid applications, *Renewable and Sustainable Energy Reviews* 44 (2015) 751–766. <https://doi.org/10.1016/j.rser.2015.01.016>.
- [22] M.W. Khan, J. Wang, M. Ma, L. Xiong, P. Li, F. Wu, Optimal energy management and control aspects of distributed microgrid using multi-agent systems, *Sustain Cities Soc* 44 (2019) 855–870. <https://doi.org/10.1016/j.scs.2018.11.009>.
- [23] A. Hussain, V.-H. Bui, H.-M. Kim, Microgrids as a resilience resource and strategies used by microgrids for enhancing resilience, *Appl Energy* 240 (2019) 56–72. <https://doi.org/10.1016/j.apenergy.2019.02.055>.
- [24] M.F. Rabbi, Cross-framework hybrid artificial intelligence for high-penetration renewable energy integration: Multi-regional forecasting and adaptive control, *Appl Energy* 401 (2025) 126834. <https://doi.org/10.1016/J.APENERGY.2025.126834>.