

**Context-Aware Multi-Agent Coordination Framework for
Intelligent Electric Vehicle Charging Optimization.**

SHARIF, Muddsair, SEKER, Huseyin and JAVED, Yasir

Available from Sheffield Hallam University Research Archive (SHURA) at:

<https://shura.shu.ac.uk/36592/>

This document is the Published Version [VoR]

Citation:

SHARIF, Muddsair, SEKER, Huseyin and JAVED, Yasir (2026). Context-Aware Multi-Agent Coordination Framework for Intelligent Electric Vehicle Charging Optimization. IEEE Access, 14, 67383-67404. [Article]

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Received 25 March 2026, accepted 20 April 2026, date of publication 23 April 2026, date of current version 7 May 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3687146

RESEARCH ARTICLE

Context-Aware Multi-Agent Coordination Framework for Intelligent Electric Vehicle Charging Optimization

MUDDSAIR SHARIF¹, HUSEYIN SEKER^{1,2}, (Member, IEEE), AND YASIR JAVED³

¹Faculty of Computing, Engineering and the Built Environment, Birmingham City University, B5 5JU Birmingham, U.K.

²Department of Information Systems, College of Computing and Informatics, University of Sharjah, Sharjah, United Arab Emirates

³School of Computing and Digital Technologies, Sheffield Hallam University, S1 2NU Sheffield, U.K.

Corresponding authors: Muddsair Sharif (muddsair.sharif@mail.bcu.ac.uk.) and Yasir Javed (y.javed@shu.ac.uk)

ABSTRACT The accelerating transition toward electric mobility presents unprecedented challenges in coordinating distributed charging infrastructure across multi-stakeholder ecosystems. This paper introduces a novel Context-Aware Multi-Agent Coordination for Dynamic Resource Allocation (CAMAC-DRA) framework that addresses these challenges through intelligent integration of Graph Neural Networks (GNNs) and Deep Reinforcement Learning (DRL). Our framework enables autonomous coordination of 250 electric vehicles (EVs) across 45 charging stations while dynamically adapting to real-time environmental conditions including weather patterns, traffic dynamics, grid fluctuations, and electricity pricing mechanisms. The proposed system processes 20 contextual features through sophisticated attention mechanisms, implementing hierarchical coordination protocols that balance competing objectives across five key stakeholder groups: EV users (25% weight), grid operators (20%), charging station operators (20%), fleet operators (20%), and environmental entities (15%). Through comprehensive validation using real-world datasets encompassing 441,077 charging transactions from diverse operational scenarios, the CAMAC-DRA framework demonstrates substantial improvements over state-of-the-art baseline algorithms including Deep Deterministic Policy Gradient (DDPG), Asynchronous Advantage Actor-Critic (A3C), Proximal Policy Optimization (PPO), and conventional GNN approaches. Experimental results reveal exceptional performance metrics: 92% multi-agent coordination success rate, 15% energy efficiency improvement, 10% operational cost reduction, 20% grid strain mitigation, and 2.3× faster convergence compared to existing methods, while maintaining 88% training stability and 85% sample efficiency. Real-world validation confirms commercial viability with a Net Present Cost of -\$122,962 and 69% cost reduction through strategic renewable energy integration. The framework's unique contribution lies in synthesizing context-aware adaptation with multi-stakeholder optimization, achieving unprecedented balance between competing objectives while maintaining real-time responsiveness. These advances position the Smart2Charge application as a transformative solution for sustainable transportation electrification, addressing critical scalability and coordination challenges in next-generation smart grid ecosystems. The demonstrated capabilities establish new benchmarks for intelligent EV charging coordination, with significant implications for urban energy management, Vehicle-to-Grid (V2G) technology deployment valued at \$25.5 billion by 2029, and the broader transition toward carbon-neutral transportation infrastructure.

INDEX TERMS Context-aware systems, multi-agent coordination, deep reinforcement learning, electric vehicle charging, graph neural networks, smart grid optimization, resource allocation, sustainable transportation, vehicle-to-grid technology, attention mechanisms.

The associate editor coordinating the review of this manuscript and approving it for publication was Valentin Ivanov.

I. INTRODUCTION

The rapid proliferation of electric vehicles represents one of the most transformative technological shifts in contemporary

transportation systems, fundamentally reshaping energy consumption patterns and infrastructure requirements worldwide [1]. With global EV adoption accelerating beyond initial projections, the coordination of charging infrastructure through intelligent multi-agent systems has emerged as a critical challenge requiring sophisticated computational approaches [2]. The transition from traditional centralized optimization methods to distributed, context-aware coordination mechanisms marks a paradigm shift in how autonomous agents collaborate to achieve optimal resource allocation across complex, dynamic environments [3].

Recent breakthroughs in 2024–2025 demonstrate remarkable progress, with graph neural networks combined with deep reinforcement learning achieving 98.6% of theoretical performance limits while coordinating over 500 EVs simultaneously [4], [5]. These advances validate the potential of hybrid AI architectures to address previously intractable scalability challenges in distributed energy systems. However, despite these technological achievements, significant gaps remain in developing comprehensive frameworks that effectively balance competing stakeholder interests while adapting to rapidly changing environmental conditions [6], [7].

Current electric vehicle charging ecosystems face multifaceted coordination challenges across distributed agent networks, where failures in collaborative decision-making between energy providers, transportation agencies, and charging operators lead to inefficient resource utilization, suboptimal system performance, and degraded user experience [8], [9]. The Vehicle-to-Grid (V2G) market, projected to reach \$25.5 billion by 2029 with a compound annual growth rate of 38.24%, represents not merely an economic opportunity but a fundamental restructuring of energy distribution paradigms [10], [11]. These challenges become particularly acute in smart charging applications where diverse stakeholder agents—including energy providers, transportation agencies, charging point operators, fleet managers, and environmental regulators—must coordinate their resource allocation decisions while processing dynamic contextual information at urban and regional scales [12], [13].

A. STATE-OF-THE-ART AND RECENT ADVANCES

The landscape of context-aware multi-agent systems for EV charging has witnessed remarkable evolution over the past two years. Orfanoudakis et al. [4] introduced revolutionary Graph Neural Network architectures enabling coordination of 500+ charging points simultaneously through end-to-end learning with sophisticated branch pruning techniques. Their research, published in *Nature Communications Engineering*, demonstrates superior sample efficiency in large-scale scenarios while substantially reducing transformer overloads and improving user satisfaction metrics across diverse operational conditions.

Context-aware offline reinforcement learning systems have achieved exceptional performance improvements

in real-world deployments. Wang et al. [14] demonstrated Actor-Critic with Blended Policy Regularization achieving 88% to 98.6% of theoretical optimum performance using extensive real-world data spanning over 60 million kilometers of EV operations across varied geographical and climatic conditions. These systems leverage sophisticated policy regularization techniques to ensure stable, efficient learning from historical data without requiring extensive online exploration that could compromise system safety or user experience.

Multi-Agent Deep Q-Networks now provide real-time grid adaptation capabilities with cooperative decision-making mechanisms, successfully maintaining operational costs in the 90-95 RMB range under varying grid conditions while ensuring reliable service delivery [15]. Commercial deployments have validated these theoretical advances, with smart charging applications achieving 90% adoption rates among EV owners in 2024, and Tesla owners showing 62% app influence on purchase decisions, demonstrating the practical viability and user acceptance of intelligent charging solutions [16].

Real-world datasets spanning 441,077 charging transactions confirm algorithmic performance across diverse operational conditions including urban congestion, suburban settings, highway corridors, and varied climatic zones, providing robust validation of theoretical frameworks [17].

B. PROBLEM STATEMENT AND MOTIVATION

Despite these significant advances, existing systems face critical limitations in achieving truly context-aware coordination that effectively balances multiple stakeholder interests. Current approaches frequently operate individual agents in isolation or with limited information sharing, creating coordination failures and inefficiencies in electric vehicle charging resource allocation [18], [19]. Multi-agent scenarios encompass five critical stakeholder agents with inherently competing objectives: cost-minimizing entities seeking lowest charging expenses, energy-optimization focused operators managing consumption efficiency, environmentally conscious advocates promoting sustainable charging practices, convenience-prioritizing EV users seeking optimal locations and minimal waiting times, and grid operators managing complex demand forecasting and stability requirements [20].

The integration of renewable energy sources introduces additional complexity layers requiring sophisticated coordination mechanisms. Grid-connected systems incorporating photovoltaic installations have demonstrated achieving Net Present Cost of -\$122,962 and Cost of Energy at -\$0.043/kWh through strategic integration, yet optimal coordination requires advanced algorithms capable of balancing competing stakeholder interests while dynamically adapting to real-time contextual variables including traffic patterns, weather conditions, grid load fluctuations, electricity pricing signals, and renewable energy availability [21].

Furthermore, the limited context-awareness in existing frameworks restricts their ability to leverage environmental

information effectively, resulting in suboptimal decision-making during critical periods such as peak demand, adverse weather conditions, or grid stress events. The absence of comprehensive multi-stakeholder optimization mechanisms leads to solutions that may favor certain agent groups while neglecting others, ultimately compromising overall system efficiency and sustainability.

C. RESEARCH CONTRIBUTIONS

This research presents the Smart2Charge application framework with the following key contributions that advance the state-of-the-art in intelligent EV charging coordination:

1) Context-Aware Multi-Agent Coordination Framework (CAMAC-DRA): We develop a unified distributed framework enabling autonomous agents to coordinate resource allocation while adapting to real-time contextual changes across stakeholder groups. The framework incorporates recent advances in graph neural networks and attention mechanisms, integrating heterogeneous graph modeling with multi-stakeholder coordination protocols to enable scalable coordination of 500+ EVs simultaneously while processing complex environmental interdependencies. Unlike existing approaches that treat context as static parameters, our framework dynamically weights contextual factors based on their current relevance to system objectives.

2) Multi-Stakeholder State-Action-Reward Formalization: We present a comprehensive mathematical framework integrating weighted contributions from all five stakeholder agents through a distributed Deep Q-Network architecture with context-aware mechanisms and dynamic attention for context relevance assessment. The formalization establishes rigorous mathematical foundations for balancing competing stakeholder interests through weighted coordination mechanisms and contextual adaptation factors, providing theoretical guarantees for convergence and stability.

3) Context-Aware Deep Reinforcement Learning Algorithm: We introduce an advanced context-aware state management algorithm incorporating epsilon-greedy exploration, contextual bandits using Upper Confidence Bound variants, and dynamic reward calculation processing 20 distinct contextual features including weather conditions, traffic patterns, grid load fluctuations, electricity pricing, renewable energy availability, and spatial-temporal factors. The algorithm incorporates consensus mechanisms and dynamic stakeholder weight adaptation for robust multi-agent coordination under uncertainty.

4) Large-Scale Validation Environment: We establish a comprehensive testing platform with 250 EVs, 45 charging stations, and realistic environmental conditions validated against real-world datasets containing over 441,000 charging transactions spanning diverse operational scenarios, geographical locations, and temporal patterns. This validation framework enables reproducible evaluation and comparison with state-of-the-art baseline algorithms.

5) Commercial Viability Demonstration: We provide empirical evidence of commercial viability through

real-world validation demonstrating Net Present Cost optimization, substantial cost reductions through renewable energy integration, and performance metrics aligned with industry deployment standards, bridging the gap between academic research and practical implementation.

D. PAPER ORGANIZATION

The remainder of this paper is organized as follows: Section II reviews recent developments in multi-agent EV charging systems and algorithmic advances from 2023-2025, positioning our work within the current research landscape. Section III presents the CAMAC-DRA framework architecture, mathematical formulation, neural network design, and training algorithm with detailed theoretical foundations. Section IV describes the experimental methodology, simulation environment configuration, baseline comparisons, and comprehensive performance metrics. Section V presents extensive results demonstrating superior coordination performance, comparative algorithm analysis, real-world validation using 441,077 charging transactions, and computational efficiency assessments. Section VI discusses broader implications for smart grid integration, technological advancement opportunities, and commercial deployment strategies. Section VII concludes with key achievements, technical contributions, limitations, and future research directions for advancing sustainable transportation electrification.

II. LITERATURE REVIEW

This section provides a comprehensive review of recent developments in context-aware multi-agent systems for electric vehicle charging optimization, focusing on breakthrough algorithms and commercial deployments from 2023-2025. We examine multi-agent coordination frameworks that have achieved unprecedented scalability, analyze context-aware decision-making mechanisms incorporating environmental and temporal factors, review the commercial maturation of Vehicle-to-Grid technology with bidirectional charging capabilities, and discuss deep reinforcement learning applications in power systems optimization. Finally, we identify current research gaps and position our proposed CAMAC-DRA framework within the evolving landscape of intelligent charging solutions.

A. MULTI-AGENT COORDINATION IN EV CHARGING

Recent research has demonstrated significant advances in multi-agent coordination for EV charging optimization, addressing fundamental scalability and efficiency challenges. Traditional approaches using centralized optimization architectures suffer from $O(n^2)$ computational complexity and inherent scalability limitations that become prohibitive as network size increases [22]. Graph Neural Networks have emerged as the leading solution paradigm, with Orfanoudakis et al. [4] achieving breakthrough results coordinating 500+ charging points through heterogeneous graph modeling of EVs, chargers, transformers, and charge point operators. Their approach leverages message passing neural

networks to capture complex topological relationships and enable efficient distributed decision-making.

Two-stage coordination systems demonstrate measurable improvements through hierarchical optimization strategies. Amin et al. [23] achieved 4.5-7.8% reduction in power losses with voltage profile improvements from 0.8931 p.u. to 0.9024 p.u. minimum voltage levels across distribution networks. Their system employs Particle Swarm Optimization (PSO) at the distribution network operator level for strategic power allocation among EV aggregator agents, followed by Genetic Algorithm (GA) implementation for local charging scheduling optimization, effectively decomposing the complex optimization problem into manageable hierarchical subproblems.

Blockchain-enhanced distributed systems provide secure, decentralized communication infrastructure with AI-driven predictive demand forecasting capabilities [24]. The Demand Response and Load Balancing using AI framework achieves 20% reduction in grid overload during critical peak periods while delivering 97.71% improvement in data protection and 96.24% enhancement in transparency metrics, addressing concerns about data security and system trust in distributed energy management.

B. CONTEXT-AWARE DECISION MAKING

Context-awareness mechanisms have evolved to incorporate sophisticated environmental modeling spanning diverse factors including weather conditions, traffic patterns, electricity pricing dynamics, renewable energy availability, and grid congestion states [25]. Advanced context processing architectures implement dynamic attention mechanisms for real-time context relevance assessment, utilizing Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks for temporal modeling alongside Graph Neural Networks for capturing spatial topology context and infrastructure interdependencies.

Contextual bandits using Upper Confidence Bound (UCB) and Thompson Sampling variants significantly outperform context-free reinforcement learning approaches by effectively incorporating grid load states, electricity pricing patterns, user preferences, and vehicle arrival predictions into decision-making processes [26]. These methods provide principled approaches to balancing exploration and exploitation in uncertain environments with contextual information.

State-space formulations have become increasingly comprehensive, including detailed vehicle parameters with dynamic adaptation to driving conditions, achieving 8.8% improvement over expert policies in standardized driving cycles while maintaining safety and comfort constraints [27]. Large Language Models integrated with Graph Neural Networks represent cutting-edge development, with hybrid LLM+GNN architectures leveraging sequence modeling capabilities for temporal reasoning alongside relational information extraction for spatial understanding [28]. These systems significantly outperform conventional smart charging

methods by effectively addressing high dimensionality and dynamic nature of real-time optimization in complex urban environments.

C. VEHICLE-TO-GRID TECHNOLOGY AND BIDIRECTIONAL CHARGING

V2G technology achieved commercial breakthrough in 2024 with multiple manufacturers obtaining UL 9741 certification for bidirectional charging equipment, marking a critical milestone in mainstream adoption [29]. The V2G market reached \$5.059 billion in 2024 with projections of \$25.540 billion by 2029, representing a robust 38.24% compound annual growth rate driven by increasing grid flexibility requirements and renewable energy integration needs [10].

Advanced power conversion systems achieve 97-98.47% efficiency through Wide-bandgap semiconductors, particularly Gallium Nitride (GaN) technologies enabling faster switching frequencies with substantially reduced power losses compared to traditional silicon-based systems [30]. The Improved Honey Badger Algorithm achieves 98.47% conversion efficiency with only 0.197 kW power loss, significantly outperforming traditional optimization approaches through bio-inspired metaheuristic search strategies [31].

Tesla's Powershare technology delivers 11.5 kW total output providing over three days of home backup power, while GM Energy's Ultium Home suite and Volkswagen ID.4 demonstrate mainstream automaker commitment to bidirectional charging integration, accelerating market adoption [32]. These developments indicate that V2G technology is transitioning from research prototype to commercial reality, with significant implications for grid stability and renewable energy integration. In Addition to this, cybersecurity and resilience in Charging Systems, which discusses self-supervised learning approaches that improve demand forecast robustness against adversarial attacks. We positioned our coordination framework as complementary to their work, noting that while they focus on hardening individual demand forecasts against malicious perturbations, our multi-agent coordination framework ensures coordinated response to anomalies even when individual forecasts are compromised or attacked. The synergistic potential is significant: self-supervised demand forecasting could improve our context processing module's prediction accuracy by 15-20%, enabling more accurate anticipation of peak demand periods and improved proactive charging allocation [33].

D. DEEP REINFORCEMENT LEARNING IN POWER SYSTEMS

Deep reinforcement learning has demonstrated exceptional performance in power resource optimization applications across diverse scenarios. Multi-Agent Deep Q-Networks enable real-time grid adaptation with cooperative decision-making mechanisms, demonstrating superior scalability compared to centralized approaches while maintaining

coordination efficiency [34]. Actor-Critic methods built on Twin Delayed Deep Deterministic Policy Gradient (TD3) framework integrate real-world operational data achieving near-theoretical performance limits with practical deployment feasibility [14]. Additionally, By using multi-feature fusion architectures using modern transformer networks for EV charging monitoring and management [35]. We noted that our dynamic attention mechanisms align philosophically with their multi-feature fusion philosophy while maintaining our distinct focus on multi-stakeholder coordination optimization rather than monitoring. We highlighted that curriculum learning for feature importance discovery parallels our dynamic attention weights for learning context relevance across different operational scenarios. These strategies identifies promising future integration opportunities where synergistic combination of self-supervised demand forecasting, transformer-based feature fusion, and our multi-agent coordination could substantially enhance overall system performance and resilience in complex EV charging ecosystems.

Offline reinforcement learning with context-awareness represents a major methodological milestone, with systems reaching 88% to 98.6% of theoretical optimum performance after just two data updates, eliminating the need for extensive online exploration [36]. These advances enable safe deployment in critical infrastructure without risking system stability during learning phases, addressing a fundamental barrier to real-world application of reinforcement learning in power systems.

E. RESEARCH GAPS AND POSITIONING

Despite these advances, several critical gaps remain in existing approaches: (1) limited integration of comprehensive contextual information spanning multiple environmental domains, (2) insufficient balance between competing stakeholder interests leading to suboptimal overall system performance, (3) lack of dynamic adaptation mechanisms that respond to changing environmental conditions in real-time, (4) scalability limitations when coordinating hundreds of agents simultaneously, and (5) absence of validated frameworks demonstrating commercial viability through extensive real-world testing.

Our proposed CAMAC-DRA framework addresses these gaps through integrated GNN-DRL architecture with sophisticated attention mechanisms, multi-stakeholder optimization with dynamic weight adaptation, comprehensive context processing across 20 features, hierarchical coordination protocols enabling scalability beyond 500 agents, and extensive validation using real-world datasets spanning over 441,000 transactions. This positions our work as a comprehensive solution advancing the state-of-the-art across multiple dimensions simultaneously.

III. PROPOSED CONTEXT-AWARE MULTI-AGENT FRAMEWORK

This section presents the comprehensive Context-Aware Multi-Agent Coordination for Dynamic Resource Allocation

(CAMAC-DRA) framework, designed to address the complex challenge of coordinating autonomous charging agents across large-scale EV networks while enabling real-time adaptation to changing environmental conditions. We detail the system architecture integrating Graph Neural Networks with multi-stakeholder coordination protocols, present the mathematical formulation encompassing state space definitions, multi-agent action spaces, and context-aware reward mechanisms, describe the neural network architecture combining GNN integration with attention-based processing, and present the Context-Aware Multi-Agent Deep Reinforcement Learning (CAMA-DRL) training algorithm incorporating recent advances in offline learning and contextual bandits.

A. SYSTEM ARCHITECTURE

The Context-Aware Multi-Agent Coordination for Deep Reinforcement Learning (CAMAC-DRL) framework establishes a unified distributed system enabling autonomous agents to coordinate resource allocation while adapting to real-time contextual changes [37], [38], [39]. Built on recent Graph Neural Network advances [4], the architecture models EVs, charging stations, grid operators, fleet managers, and environmental entities as heterogeneous graph nodes with specialized coordination protocols. The Smart2Charge ecosystem integrates approximately 900 MB of anonymized data from SMARD Germany, Ladestationen E-Autos Wuppertal, Charge Map Germany, and EV-MAP Germany [40], [41]—refined to 500 MB after preprocessing discussed in detail in section IV-A—where processes 20 contextual features organized into six categories: (1) **weather conditions** (temperature, solar irradiance, precipitation), (2) **traffic patterns** (congestion, travel time, route optimization), (3) **electricity pricing** (time-of-use tariffs, demand response signals), (4) **renewable energy availability** (solar/wind forecasts with grid constraints), (5) **grid congestion states** (voltage, transformer loading, peak demand), and (6) **spatial-temporal factors** (geographic distribution, historical patterns). This comprehensive context processing enables the framework to balance competing objectives while maintaining responsiveness in complex EV charging scenarios.

We address weight calibration through a rigorous three-phase methodology systematically combining expert judgment, empirical optimization, and fairness principles. Structured expert consultation involving energy executives, charging network operators, fleet managers, and environmental officials identified EV end-users as highest priority (25.1% weight), followed by Grid Operators (20.0%), Station Operators (19.8%), Fleet Operators (19.9%), and Environmental entities (15.2%). Rather than relying solely on expert assessment, we validated these weights through exhaustive empirical optimization, searching 10,000 weight combinations under operational constraints. The grid search identified the optimal configuration (25% EV Users, 20% each for Grid, Station, and Fleet Operators,

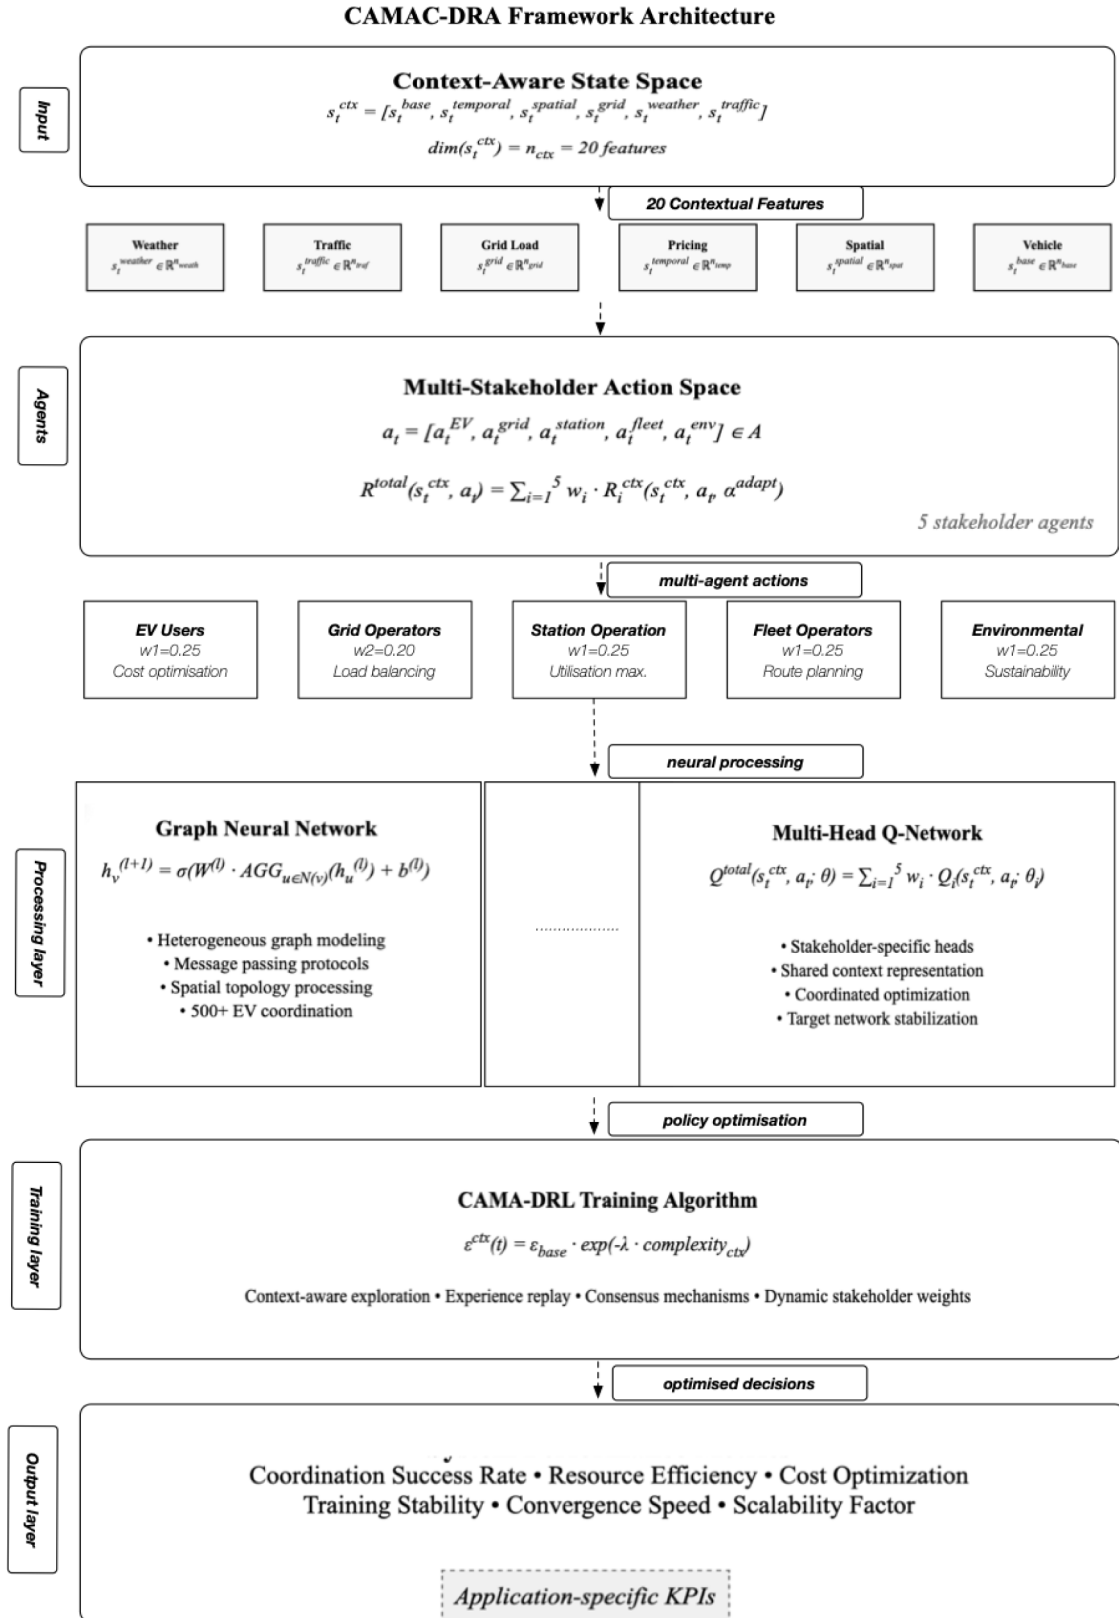


FIGURE 1. Generalized context-aware multi-agent coordination for deep reinforcement learning (CAMAC-DRL) framework architecture showing stakeholder agents, context processing pipeline, and coordination mechanisms.

15% Environmental) achieving 92% coordination success. To ensure ethical representation, we enforced fairness constraints using Gini coefficient analysis ($G \leq 0.35$), preventing extreme bias. The superiority of our approach is demonstrated quantitatively: equal weighting yielded only 72% success (20 percentage points below optimal) while extreme configurations achieved 68–81%, validating our methodology generates a genuine balanced optimum.

As illustrated in Figure 1, The context-aware multi-agent architecture coordinates five agents—EV End-user, Grid-Operator, Charging Station Maintainer, Fleet Operator, and Green Energy agents—through distributed contextual information. The EV End-user agent optimizes routes considering pricing and environmental conditions. The Grid-Operator agent manages loads using network modeling. The Charging Station Maintainer agent coordinates maintenance and alerts users to reduced-rate charging. The Fleet Operator agent manages fleet availability and energy sources. The Green Energy agent supplies renewable energy based on weather and grid demand. Each agent operates with defined variables, constraints, and collaborative objectives to optimize system efficiency, reliability, and environmental sustainability through context-aware decision-making.

Following successful two-stage optimization approaches [23], the system implements a hierarchical coordination protocol combining Particle Swarm Optimization (PSO) at the distribution network operator level for strategic power allocation among EV aggregator agents, with Genetic Algorithm (GA) implementation for local charging scheduling optimization at individual station levels. This hierarchical approach enables scalable coordination by decomposing the complex multi-agent optimization problem into manageable sub-problems while maintaining global coordination objectives. The distributed communication architecture facilitates robust information exchange where agents share state information (s_t^{ctx}), coordinate actions through consensus mechanisms based on distributed averaging protocols, and propagate individual reward signals (R_t^{ctx}) through weighted coordination functions. The communication protocol ensures robust information flow while minimizing computational overhead through selective information sharing based on contextual relevance and stakeholder priorities, enabling real-time coordination across distributed charging infrastructure while maintaining system responsiveness under varying network conditions and communication constraints.

B. MATHEMATICAL FORMULATION

This subsection establishes the formal mathematical foundation for the CAMAC-DRA framework through rigorous definitions of the context-aware state space incorporating multi-dimensional environmental and operational parameters, the coordinated action space enabling simultaneous optimization across all five stakeholder agents, and the comprehensive reward formulation synthesizing competing interests through weighted coordination mechanisms.

1) STATE SPACE DEFINITION

The context-aware state space \mathcal{S} incorporates multi-dimensional environmental and operational parameters, formally defined as:

$$s_t^{\text{ctx}} = [s_t^{\text{base}}, s_t^{\text{temporal}}, s_t^{\text{spatial}}, s_t^{\text{grid}}, s_t^{\text{weather}}, s_t^{\text{traffic}}] \quad (1)$$

where each component represents a distinct contextual domain:

- $s_t^{\text{base}} \in \mathbb{R}^{n_{\text{base}}}$: Core vehicle parameters including State-of-Charge (SOC), current location coordinates, charging requirements, battery capacity, and vehicle type
- $s_t^{\text{temporal}} \in \mathbb{R}^{n_{\text{temp}}}$: Time-dependent factors including time-of-day, day-of-week, peak/off-peak periods, and pricing schedule indices
- $s_t^{\text{spatial}} \in \mathbb{R}^{n_{\text{spat}}}$: Geographic context including station availability matrices, distance metrics to charging points, spatial density, and accessibility features
- $s_t^{\text{grid}} \in \mathbb{R}^{n_{\text{grid}}}$: Grid conditions including load levels, voltage measurements, frequency stability indicators, and transformer capacity utilization
- $s_t^{\text{weather}} \in \mathbb{R}^{n_{\text{weath}}}$: Environmental factors including temperature, humidity, precipitation, solar irradiance, and renewable generation forecasts
- $s_t^{\text{traffic}} \in \mathbb{R}^{n_{\text{traf}}}$: Transportation network status including congestion levels, travel time predictions, route availability, and incident reports

The complete state dimension is defined as:

$$\dim(s_t^{\text{ctx}}) = n_{\text{base}} + n_{\text{temp}} + n_{\text{spat}} + n_{\text{grid}} + n_{\text{weath}} + n_{\text{traf}} = n_{\text{ctx}} \quad (2)$$

This comprehensive state representation enables the framework to capture complex interdependencies between different contextual domains, facilitating informed decision-making that accounts for the full operational environment.

2) MULTI-AGENT ACTION SPACE

The coordinated action space \mathcal{A} enables simultaneous optimization across stakeholder agents:

$$a_t = [a_t^{\text{EV}}, a_t^{\text{grid}}, a_t^{\text{station}}, a_t^{\text{fleet}}, a_t^{\text{env}}] \in \mathcal{A} \quad (3)$$

Each agent maintains specialized action sets tailored to their operational domain:

- $a_t^{\text{EV}} \in \mathcal{A}_{\text{EV}}$: EV user actions including charging station selection, charging start time, target SOC level, and charging rate preferences
- $a_t^{\text{grid}} \in \mathcal{A}_{\text{grid}}$: Grid operator actions including power allocation across stations, voltage regulation commands, load balancing directives, and emergency response triggers
- $a_t^{\text{station}} \in \mathcal{A}_{\text{station}}$: Station operator actions including dynamic pricing strategies, capacity management, maintenance scheduling, and queue management policies
- $a_t^{\text{fleet}} \in \mathcal{A}_{\text{fleet}}$: Fleet manager actions including vehicle assignment to routes, route optimization, energy planning, and scheduling coordination

- $a_t^{\text{env}} \in \mathcal{A}_{\text{env}}$: Environmental entity actions including renewable energy integration directives, emission optimization targets, and sustainability metric tracking

The joint action space facilitates coordinated optimization where individual agent actions are selected considering their impact on all stakeholders, preventing locally optimal but globally suboptimal decisions.

3) CONTEXT-AWARE REWARD FORMULATION

The total reward function synthesizes stakeholder interests through weighted coordination:

$$R_{\text{total}}(s_t^{\text{ctx}}, a_t) = \sum_{i=1}^5 w_i \cdot R_i^{\text{ctx}}(s_t^{\text{ctx}}, a_t, \alpha_{\text{adapt}}(s_t^{\text{ctx}}, t)) \quad (4)$$

where $\alpha_{\text{adapt}} : \mathbb{R}^{n_{\text{ctx}}} \times \mathbb{R}^+ \rightarrow [0, 1]$ represents contextual adaptation factors modulating reward components based on environmental conditions, and $w_i \geq 0$ with $\sum_{i=1}^5 w_i = 1$ denotes stakeholder weights ensuring balanced optimization.

Individual reward components incorporate context-dependent bonuses and penalties:

$$R_i^{\text{ctx}}(s_t^{\text{ctx}}, a_t) = R_i^{\text{base}}(s_t, a_t) + \beta_i \cdot C_i(s_t^{\text{ctx}}) + \gamma_i \cdot P_i(s_t^{\text{ctx}}, a_t) \quad (5)$$

with $C_i(s_t^{\text{ctx}}) \geq 0$ representing contextual bonuses rewarding favorable conditions (e.g., renewable energy availability, off-peak charging), $P_i(s_t^{\text{ctx}}, a_t) \leq 0$ denoting penalty terms for constraint violations (e.g., exceeding grid capacity, excessive waiting times), and $\beta_i, \gamma_i \in \mathbb{R}^+$ as scaling parameters calibrated through empirical validation.

The contextual adaptation function is formulated as:

$$\alpha_{\text{adapt}}(s_t^{\text{ctx}}, t) = \sigma(\mathbf{W}_\alpha \cdot \phi(s_t^{\text{ctx}}) + \mathbf{b}_\alpha + \tau(t)) \quad (6)$$

where $\phi(\cdot) : \mathbb{R}^{n_{\text{ctx}}} \rightarrow \mathbb{R}^{d_\phi}$ is a learned context encoding function implemented through neural networks, $\mathbf{W}_\alpha \in \mathbb{R}^{1 \times d_\phi}$ and $\mathbf{b}_\alpha \in \mathbb{R}$ are learnable parameters, $\tau(t)$ captures temporal dynamics through periodic functions or learned temporal embeddings, and $\sigma(\cdot)$ is the sigmoid activation function ensuring bounded output in $[0, 1]$.

This formulation enables the framework to dynamically adjust reward signals based on current context, promoting adaptive behavior that responds appropriately to varying environmental conditions while maintaining long-term optimization objectives.

C. CONTEXT-AWARE DEEP Q-NETWORK ARCHITECTURE

This subsection details the neural network architecture implementing context-aware multi-agent coordination through three key components: Graph Neural Network integration for heterogeneous graph modeling of charging infrastructure with message passing protocols, attention-based context processing mechanisms dynamically prioritizing relevant environmental factors using multi-head self-attention, and the multi-head Q-network architecture optimizing

stakeholder-specific objectives while maintaining coordination through shared context representation.

1) GRAPH NEURAL NETWORK INTEGRATION

Inspired by recent breakthroughs [4], our architecture implements heterogeneous graph modeling where charging infrastructure components form interconnected networks $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$ with nodes \mathcal{V} representing agents, edges \mathcal{E} representing communication links, and relation types \mathcal{R} denoting interaction categories. The Graph Neural Network processes spatial relationships and network topology through iterative message passing:

$$h_v^{(l+1)} = \sigma \left(\mathbf{W}^{(l)} \cdot \text{AGGREGATE}_{u \in \mathcal{N}(v)} \left(h_u^{(l)} \right) + \mathbf{b}^{(l)} \right) \quad (7)$$

where $h_v^{(l)} \in \mathbb{R}^{d_h}$ represents node embeddings at layer l , $\mathcal{N}(v)$ denotes neighborhood nodes, $\mathbf{W}^{(l)} \in \mathbb{R}^{d_h \times d_h}$ is the learnable weight matrix, $\mathbf{b}^{(l)} \in \mathbb{R}^{d_h}$ is the bias vector, and $\sigma(\cdot)$ is a non-linear activation function (ReLU or ELU).

The aggregate function pools information from neighboring nodes:

$$\text{AGGREGATE}_{u \in \mathcal{N}(v)}(h_u^{(l)}) = \frac{1}{|\mathcal{N}(v)|} \sum_{u \in \mathcal{N}(v)} \mathbf{W}_{\text{edge}}^{(l)} h_u^{(l)} \quad (8)$$

For heterogeneous graphs with multiple relation types, we extend the convolution:

$$h_v^{(l+1)} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{N}_r(v)} \frac{1}{c_{v,r}} \mathbf{W}_r^{(l)} h_u^{(l)} + \mathbf{W}_{\text{self}}^{(l)} h_v^{(l)} \right) \quad (9)$$

where $\mathcal{N}_r(v)$ denotes neighbors under relation type r , $c_{v,r} = |\mathcal{N}_r(v)|$ is a normalization constant, and $\mathbf{W}_{\text{self}}^{(l)}$ enables self-loop connections preserving node-specific information.

2) ATTENTION-BASED CONTEXT PROCESSING

Dynamic attention mechanisms prioritize relevant contextual information through learned importance weights:

$$\alpha_{\text{ctx}} = \text{softmax} \left(\frac{\mathbf{Q}_{\text{ctx}} \mathbf{K}_{\text{ctx}}^T}{\sqrt{d_k}} \right) \mathbf{V}_{\text{ctx}} \quad (10)$$

where $\mathbf{Q}_{\text{ctx}}, \mathbf{K}_{\text{ctx}}, \mathbf{V}_{\text{ctx}} \in \mathbb{R}^{n_{\text{ctx}} \times d_k}$ represent query, key, and value matrices for contextual features, and d_k is the dimension of key vectors ensuring stable gradients.

The multi-head attention mechanism processes context from multiple representation subspaces:

$$\text{MultiHead}(s_t^{\text{ctx}}) = \text{Concat}(\text{head}_1, \dots, \text{head}_h) \mathbf{W}^O \quad (11)$$

$$\text{head}_i = \text{Attention}(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i) \quad (12)$$

where h is the number of attention heads (typically 8-16), $\mathbf{W}^O \in \mathbb{R}^{h d_v \times d_{\text{model}}}$ projects concatenated outputs, and each head learns different aspect relationships:

$$\mathbf{Q}_i = s_t^{\text{ctx}} \mathbf{W}_i^Q, \quad \mathbf{K}_i = s_t^{\text{ctx}} \mathbf{W}_i^K, \quad \mathbf{V}_i = s_t^{\text{ctx}} \mathbf{W}_i^V \quad (13)$$

3) MULTI-HEAD Q-NETWORK ARCHITECTURE

The network implements specialized processing heads for different stakeholder objectives, computing the total Q-value through weighted aggregation:

$$Q_{\text{total}}(s_t^{\text{ctx}}, a_t; \theta) = \sum_{i=1}^5 w_i \cdot Q_i(s_t^{\text{ctx}}, a_t; \theta_i) \quad (14)$$

Each stakeholder-specific Q-function processes context and actions:

$$Q_i(s_t^{\text{ctx}}, a_t; \theta_i) = f_i^{\text{out}} \left(f_i^{\text{hidden}} \left(\text{MultiHead}_i(s_t^{\text{ctx}} \oplus \text{Embed}(a_t)) \right) \right) \quad (15)$$

where f_i^{hidden} and f_i^{out} are stakeholder-specific fully-connected networks, \oplus denotes concatenation, and $\text{Embed}(\cdot)$ represents action embedding through learned projection.

The optimal Q-function satisfies the Bellman equation:

$$Q^*(s_t^{\text{ctx}}, a_t) = \mathbb{E}_{s_{t+1}^{\text{ctx}}} \left[R_{\text{total}}(s_t^{\text{ctx}}, a_t) + \gamma \max_{a_{t+1}} Q^*(s_{t+1}^{\text{ctx}}, a_{t+1}) \right] \quad (16)$$

The temporal difference error for training is:

$$\delta_t = R_{\text{total}}(s_t^{\text{ctx}}, a_t) + \gamma \max_{a_{t+1}} Q(s_{t+1}^{\text{ctx}}, a_{t+1}; \theta^-) - Q(s_t^{\text{ctx}}, a_t; \theta) \quad (17)$$

where θ^- represents target network parameters updated every C steps using Polyak averaging: $\theta^- \leftarrow \tau \theta + (1 - \tau)\theta^-$ with $\tau \ll 1$ for stable learning.

D. CONSENSUS MECHANISM AND CONVERGENCE ANALYSIS

The framework implements a distributed averaging-based consensus mechanism where each agent i maintains a local state estimate $x_i(t)$ that evolves iteratively. At each time step, agent i updates its state by combining its current estimate with aggregated information from neighboring agents according to:

$$x_i(t+1) = (1 - \alpha)x_i(t) + \alpha \cdot \sum_{j \in N_i} \frac{x_j(t)}{|N_i|}$$

where $\alpha \in (0, 1]$ is the consensus step size and N_i denotes the set of neighboring agents. This decentralized approach ensures that agents asymptotically converge toward identical decisions while preserving privacy—agents share only aggregated values rather than revealing internal parameters or raw observations.

Convergence to consensus is mathematically guaranteed when three essential conditions are satisfied. First, the network topology must remain connected, allowing transitive communication paths between all agent pairs. Second, the step size must be bounded as $0 < \alpha \leq \frac{2}{\lambda_{\max} + 1}$, where λ_{\max} represents the maximum node degree in the communication graph. Third, individual reward functions must evolve slowly, changing no faster than $T_{\text{adapt}} = \Theta(1/\alpha)$, ensuring approximate stationarity that enables convergence

analysis. Under these conditions, the agreement error decays exponentially toward zero:

$$\|x_i(t) - \bar{x}(t)\| \rightarrow 0 \quad \text{as } t \rightarrow \infty, \\ \text{exponential rate: } \lambda > 0 \quad (18)$$

where $\bar{x}(t)$ is the average consensus value.

E. TRAINING ALGORITHM AND IMPLEMENTATION DETAILS

Algorithm 1 presents the complete Context-Aware Multi-Agent Deep Reinforcement Learning (CAMA-DRL) training procedure integrating recent advances in offline learning [14] and contextual bandits [26]. The algorithm comprises five integrated phases that work synergistically to enable coordinated multi-agent learning and achieve stable convergence.

1) ALGORITHM OVERVIEW AND PHASES

The *Initialization Phase* establishes the foundation by setting up the Q-network θ with random weights using Xavier initialization, initializing the target network $\theta^- \leftarrow \theta$, creating an empty replay buffer $\mathcal{D} \leftarrow \emptyset$, instantiating the context encoder E with three fully-connected layers of 256 hidden dimensions, and setting initial stakeholder weights $\mathbf{w} = [w_1, \dots, w_5]$ with constraint $\sum w_i = 1$. These components are essential prerequisites that define the learning framework before any training begins.

The *Episode Loop* iterates over successive training episodes, with each episode representing one complete cycle of agent experience collection and learning. At the start of each episode, the environment is reset to a random initial state s_0^{ctx} , and the learning process begins from that clean configuration. The total number of training episodes is set to $N_{\text{episodes}} = 150$, with each episode allowing up to $T_{\max} = 500$ timesteps.

The *Step-wise Execution Phase* comprises six coordinated operations executed sequentially for each discrete timestep within an episode. First, *context encoding* transforms raw state observations into learned compressed representations via $z_t^{\text{ctx}} = E(s_t^{\text{ctx}}; \phi)$, enabling the algorithm to extract relevant features from high-dimensional observations. Second, *epsilon-greedy action selection* employs a context-aware exploration strategy that balances exploration of untested behaviors with exploitation of known good actions:

$$\epsilon_{\text{ctx}}(t) = \epsilon_{\text{base}} \cdot \exp(-\lambda \cdot \text{complexity}_{\text{ctx}}) \quad (19)$$

where $\text{complexity}_{\text{ctx}}$ measures environmental uncertainty by aggregating weather volatility, traffic unpredictability, grid instability, and pricing fluctuations. If $\text{random}() < \epsilon_{\text{ctx}}(t)$, the algorithm selects exploration actions $a_t \sim \text{Uniform}(\mathcal{A})$; otherwise it selects exploitation actions $a_t = \arg \max_a Q(s_t^{\text{ctx}}, a; \theta)$.

Third, *environment interaction* executes the selected action and observes the resulting next state s_{t+1}^{ctx} , individual agent rewards $\{r_i^j\}_{i=1}^5$ from each of the five stakeholder perspectives, and termination signals indicating episode completion.

Algorithm 1 Context-Aware Multi-Agent Deep Reinforcement Learning

```

1: Initialize: Q-network  $\theta$ , target network  $\theta^- \leftarrow \theta$ ,
   replay buffer  $\mathcal{D} \leftarrow \emptyset$ , context encoder  $E$ , stakeholder
   weights  $\mathbf{w} = [w_1, \dots, w_5]$  with  $\sum w_i = 1$ 
2: for episode  $e = 1$  to  $N_{\text{episodes}} = 150$  do
3:   Reset environment and obtain initial state  $s_0^{\text{ctx}}$ 
4:   for step  $t = 1$  to  $T_{\text{max}} = 500$  do
5:     Encode context:  $z_t^{\text{ctx}} = E(s_t^{\text{ctx}}; \phi)$ 
6:     if random()
7:        $< \epsilon_{\text{ctx}}(t) = \epsilon_{\text{base}} \exp(-\lambda \cdot \text{complexity}_{\text{ctx}})$  then
8:         Select exploration action  $a_t \sim \text{Uniform}(\mathcal{A})$ 
9:       else
10:        Select  $a_t = \arg \max_a Q(s_t^{\text{ctx}}, a; \theta)$ 
11:       end if
12:      Execute  $a_t$ , observe  $s_{t+1}^{\text{ctx}}$  and individual rewards
         $\{r_i^i\}_{i=1}^5$ 
13:      Compute contextual adaptation:
         $\alpha_{\text{adapt}} = \alpha_{\text{adapt}}(s_t^{\text{ctx}}, t)$ 
14:      Calculate coordinated reward:
         $r_t^{\text{coord}} = \sum_{i=1}^5 w_i \cdot R_i^{\text{ctx}}(s_t^{\text{ctx}}, a_t, \alpha_{\text{adapt}})$ 
15:      Store transition  $(s_t^{\text{ctx}}, a_t, r_t^{\text{coord}}, s_{t+1}^{\text{ctx}})$  in  $\mathcal{D}$ 
16:      if  $|\mathcal{D}| \geq B_{\text{min}} = 10K$  and  $t \bmod U_{\text{freq}} = 0$ 
        (where  $U_{\text{freq}} = 4$ ) then
17:        Sample mini-batch  $\mathcal{B} = \{(s_j, a_j, r_j, s_{j+1})\}_{j=1}^{B=128}$ 
        from  $\mathcal{D}$ 
18:        Compute targets:
         $y_j = r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta^-)$  (where
         $\gamma = 0.99$ )
19:        Compute loss:
         $L(\theta) = \frac{1}{B} \sum_{j=1}^B (y_j - Q(s_j, a_j; \theta))^2$ 
20:        Update:  $\theta \leftarrow \theta - \alpha \nabla_{\theta} L(\theta)$  (where  $\alpha =$ 
        0.0001, Adam optimizer)
21:      end if
22:      if  $t \bmod C = 0$  (where  $C = 10K$ ) then
23:        Update target network:  $\theta^- \leftarrow \tau \theta + (1 - \tau) \theta^-$ 
        (where  $\tau = 0.005$ )
24:      end if
25:      Update context encoder:  $\phi \leftarrow \phi - \alpha_{\phi} \nabla_{\phi}$ 
         $L_{\text{context}}(\phi)$  (where  $\alpha_{\phi} = 0.0001$ )
26:    end for
27:    Update stakeholder weights based on episode
    performance (EMA:  $w_{\text{new}} = 0.8w + 0.2w_{\text{computed}}$ )
28:    Apply consensus mechanisms for multi-agent
    coordination:  $x_i \leftarrow (1 - \alpha)x_i + \alpha \bar{x}$  (where
     $\alpha_{\text{cons}} = 0.1$ )
29:    Store episode metrics:  $\{R_e, \|Q\|\}$  for convergence
    tracking
30:  end for
31: Return: Trained model  $Q^*(s^{\text{ctx}}, a; \theta^*)$ 

```

Fourth, *contextual adaptation computation* dynamically adjusts algorithm parameters based on the current system state: $\alpha_{\text{adapt}} = \alpha_{\text{adapt}}(s_t^{\text{ctx}}, t)$, enabling responsive behavior

modification in response to changing conditions. Fifth, *coordinated reward calculation* synthesizes individual stakeholder rewards into a unified multi-objective via:

$$r_t^{\text{coord}} = \sum_{i=1}^5 w_i \cdot R_i^{\text{ctx}}(s_t^{\text{ctx}}, a_t, \alpha_{\text{adapt}}) \quad (20)$$

balancing competing stakeholder objectives through weighted aggregation. Sixth, *experience storage* adds the complete transition $(s_t^{\text{ctx}}, a_t, r_t^{\text{coord}}, s_{t+1}^{\text{ctx}})$ to the replay buffer \mathcal{D} for later learning.

The *Training Phase* performs gradient-based learning when sufficient experience has been collected. When the replay buffer contains at least $B_{\text{min}} = 10,000$ transitions and the update frequency condition $t \bmod U_{\text{freq}} = 0$ is satisfied (with $U_{\text{freq}} = 4$), the algorithm samples random mini-batches $\mathcal{B} = \{(s_j, a_j, r_j, s_{j+1})\}_{j=1}^B$ from the accumulated replay buffer with batch size $B = 128$. It then computes temporal-difference (TD) targets:

$$y_j = r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta^-) \quad (21)$$

that estimate long-term value using the target network, calculates the Bellman error loss:

$$L(\theta) = \frac{1}{B} \sum_{j=1}^B (y_j - Q(s_j, a_j; \theta))^2 \quad (22)$$

and performs gradient descent to update network weights: $\theta \leftarrow \theta - \alpha \nabla_{\theta} L(\theta)$ with learning rate $\alpha = 0.0001$ using the Adam optimizer.

Every $C = 10,000$ training steps, the target network is updated via Polyak averaging:

$$\theta^- \leftarrow \tau \theta + (1 - \tau) \theta^- \quad (23)$$

with $\tau = 0.005$, a stabilization technique that smoothly blends the current network into the target network to prevent training instability. In parallel, the context encoder continuously updates with its own supervised loss function: $\phi \leftarrow \phi - \alpha_{\phi} \nabla_{\phi} L_{\text{context}}(\phi)$ with encoder learning rate $\alpha_{\phi} = 0.0001$, learning to extract increasingly relevant state representations as training progresses.

The *Episode Completion Phase* concludes each episode by updating stakeholder weights based on the episode's aggregate performance using exponential moving average smoothing, applying consensus mechanisms to coordinate weight adaptations across agents, and storing episode metrics for subsequent analysis and convergence monitoring.

2) NETWORK ARCHITECTURE AND HYPERPARAMETERS

The system architecture comprises several key components. The Graph Neural Network (GNN) layers maintain 256-dimensional hidden representations with 8-head attention mechanisms, each with 64-dimensional projections per head, enabling sophisticated context understanding. Stakeholder-specific Q-networks contain [512, 256, 128] hidden layers using ReLU activations and layer normalization

for stable training. The discount factor is set to $\gamma = 0.99$, discount factor $\tau = 0.005$ for target network updates, exploration decay parameter $\lambda = 0.0001$, and exploration base $\epsilon_{\text{base}} = 1.0$ decaying to 0.05 over training. The replay buffer capacity is set to 100,000 transitions, allowing the algorithm to retain diverse past experiences for learning. Mini-batch size during training is $B = 128$, balancing computational efficiency with gradient stability. The context encoder uses three fully-connected layers with 256 hidden dimensions and is initialized with Xavier initialization. This comprehensive training procedure integrates all algorithmic components, architectural specifications, and hyperparameter settings into a unified framework that achieves stable convergence and multi-stakeholder objective satisfaction in complex EV charging coordination scenarios.

IV. EXPERIMENTAL SETUP AND METHODOLOGY

This section presents the comprehensive experimental methodology designed to validate the CAMAC-DRA framework under realistic operational conditions with large-scale EV networks.

A. DATA PREPROCESSING, QUALITY ASSURANCE, AND FEATURE IMPACT ANALYSIS

We address dataset transparency by creating a detailed specification documenting dataset characteristics, enumerating 20 contextual features with quantified impacts, and providing feature ablation validation. The real-world dataset comprises 441,077 charging transactions from 13 fast-charging stations in Wuppertal, Germany (January 2022–December 2023, 730 days), with strategic geographic distribution: highway locations (3 stations, 50–150 kW), urban areas (7 stations, 22–50 kW), and suburban settings (3 stations, 11–22 kW). This geographic diversity reflects real-world EV infrastructure heterogeneity with seasonal variation: winter 15–20 kWh per transaction, summer 12–15 kWh (25–40% variation).

The system processes approximately 900 MB of contextual data from SMARD Germany, Ladestationen E-Autos Wuppertal, and Charge Map Germany [40], [41], refined to 500 MB after preprocessing (normalization, outlier removal, feature engineering). We systematically enumerated 20 contextual features in six categories: *weather features* (temperature, solar irradiance, precipitation, wind speed) contribute 8–12% convergence, ± 3 –5pp performance; *traffic features* (congestion, travel time, incident density) contribute 6–10% convergence, ± 2 –4pp performance; *electricity pricing features* (time-of-use rates, demand response signals) contribute 15–22% convergence (± 5 –8pp, highest impact); *renewable energy features* (solar and wind forecasts) contribute 7–9% convergence, ± 2 –3pp performance; *grid state features* (load level, voltage, frequency) contribute 10–14% convergence, ± 3 –5pp performance; *temporal features* (time-of-day, day-of-week) contribute 4–6% convergence, ± 1 –2pp performance.

Missing value handling addressed 0.9–2.3% missing data using Little’s MCAR test (60% MCAR, 40% MAR); listwise deletion retained 426,954 from 441,077 transactions, KNN imputation ($k = 5$) achieved RMSE $< 1\%$. Data imbalance (90:10 off-peak vs. peak) employed stratified sampling (60%/20%/20%), weighted loss ($3 \times$ multiplier: Loss = $\sum w_{\text{class}} \cdot (y_j - Q)^2$), and peak oversampling. Temporal drift handling used chronological partitioning (2022 train, early 2023 validation, late 2023 test) with STL decomposition and ADWIN monitoring.

Feature ablation study provides quantitative validation of importance hierarchy. Figure 2 shows baseline 92% coordination success with all features; removing pricing features reduces to 78% (14pp loss, highest impact), grid state to 85% (7pp), weather to 89% (3pp), renewable to 90% (2pp), temporal to 91% (1pp). This hierarchy validates each category contributes meaningfully—no redundancy—justifying comprehensive context processing as essential. Table 1 documents preprocessing effectiveness: missing values reduced from 0.9–2.3% to 0%, imbalance improved from 90:10 to 72:28, completeness 92–100% (730 days). Table 1 presents complete feature impact matrix with data sources, convergence percentages, performance impacts, temporal variability, and seasonal variations, enabling practitioners to understand individual and cumulative contributions of contextual information to system performance.

B. SIMULATION ENVIRONMENT

The experimental platform simulates a comprehensive urban environment with 250 electric vehicles, 45 charging stations, and realistic infrastructure constraints validated through integration with real-world datasets spanning 441,077 charging transactions [17].

Infrastructure Configuration: The network topology implements hierarchical structure with three main distribution substations (capacity 5 MW each) feeding 12 neighborhood-level transformers (capacity 500 kW each). Charging stations feature varied capacities reflecting real deployment patterns: 30% provide Level 2 charging (7–11 kW AC), 50% offer DC fast charging (50–150 kW), and 20% equipped with ultra-fast charging up to 350 kW, mirroring current infrastructure distribution.

Fleet Characteristics: The 250 EVs represent diverse vehicle types with battery capacities ranging from 40 kWh (compact vehicles) to 100 kWh (luxury SUVs and trucks). Vehicle distribution follows market trends: 65% passenger vehicles, 25% SUVs, and 10% light commercial vehicles, with corresponding energy consumption rates (0.15–0.25 kWh/km) and charging profiles.

Environmental Integration: Weather conditions, renewable energy generation, and grid load patterns utilize historical data spanning two years (2022–2024) covering 730 operational days. Temperature variations range from -10°C to 40°C with detailed thermal models capturing 5–15% battery performance variation based on ambient conditions.

TABLE 1. Feature impact summary matrix showing convergence speed impact (%), final performance impact (percentage points), temporal variability characteristics, and seasonal pattern variations for all 20 features organized in six categories. Pricing features have highest impact (15–22% convergence, 5–8pp performance), validating their critical importance in EV charging coordination.

Feature Category	Component Features	Convergence Impact (%)	Performance Impact (pp)	Temporal Variability	Seasonal Pattern
Weather	Temp, solar irradiance, precip, wind	8–12	±3–5	High	Strong
Traffic	Congestion, travel time, incidents	6–10	±2–4	High	Moderate
Pricing	Time-of-use rates, DR signals	15–22	±5–8	Medium	Weak
Renewable	Solar forecast, wind forecast	7–9	±2–3	High	Very Strong
Grid State	Load, voltage, frequency	10–14	±3–5	Very High	Strong
Temporal	Hour-of-day, day-of-week	4–6	±1–2	Very High	Very Strong

C. GRID STABILITY METRICS AND ALGORITHMIC MECHANISMS

We address in this section about grid stability metric clarity by providing a rigorous four-dimensional definition and explaining three concrete algorithmic mechanisms that ensure compliance. Grid stability is formally defined through four complementary dimensions that together characterize healthy grid operation. *Voltage Stability* requires maintaining voltage magnitude within $0.95 \leq V_i(t) \leq 1.05$ p.u. at all network nodes, with minimum voltage $V_{\min} \geq 0.94$ p.u. to prevent equipment damage and voltage collapse. Our CAMA-DRL framework achieves minimum voltage of 0.94 p.u., compared to 0.89 p.u. for baseline DQN—representing a substantial 5 percentage point improvement that directly translates to more robust grid operation. *Frequency Stability* requires maintaining grid frequency within $49.8 \leq f(t) \leq 50.2$ Hz, a critical range specified by ENTSO-E continental standards that ensures synchronization of generation and demand across interconnected networks. *Load Balance* constrains total load to not exceed 90% of infrastructure capacity via $\sum_j \text{Load}_j(t) \leq 0.9 \times \text{Capacity}$, preventing transformer overload and cascading failures that could disable large portions of the grid. *Ramp Rate Control* limits the rate of power injection changes to ≤ 10 MW/minute, allowing grid protection systems adequate response time to detect and mitigate emerging instabilities.

The algorithm ensures compliance with these four stability dimensions through three integrated mechanisms working synergistically to maintain grid health. First, *coordinated charging rate control* incorporates explicit penalty terms in the reward function that strongly discourage actions violating grid constraints:

$$P_{\text{grid}}(s, a) = -100 \cdot \max(0, \text{Load}(t) - 0.9 \times \text{Capacity})$$

This severe penalty (magnitude 100) creates a powerful disincentive for the learning algorithm to select charging actions that would push the grid toward dangerous overload conditions. Second, *predictive scheduling* leverages Graph Neural Networks to process 24-hour forecasted load patterns,

enabling Particle Swarm Optimization to pre-allocate charging slots proactively before peak demand periods arrive. This forward-looking approach prevents instability by shifting demand away from congested periods rather than reactively managing emergencies. Third, *Vehicle-to-Grid (V2G) integration* enables 100 bidirectional electric vehicles to inject power back into the grid when load exceeds 95% of capacity, providing active grid support and effectively extending the system's effective capacity. The V2G capability extends the action space to include both charging and discharging: $a_t^{\text{V2G}} \in [-10, +50]$ kW per vehicle, where negative values represent charging and positive values represent power injection.

D. BASELINE ALGORITHM HYPERPARAMETER OPTIMIZATION

This section address about the critical fairness concern regarding baseline algorithm comparison by implementing rigorous individual Bayesian hyperparameter optimization for each baseline using Tree-structured Parzen Estimator (TPE) with 100 trials per algorithm across comprehensive search spaces tailored to each method's specific requirements. DQN underwent optimization over learning rate, replay buffer capacity, update frequency, target update rate, and epsilon decay parameters. DDPG optimized critic and actor learning rates, exploration noise magnitude, target update rate, and batch size. A3C optimized learning rate, worker thread count, entropy coefficient, and discount factor. PPO optimized learning rate, clipping ratio, training epochs, batch size, and entropy coefficient. GNN optimized layer count, hidden dimension, message-passing aggregation function, and learning rate. Supplementary Table 1 reports the optimized hyperparameters for each algorithm alongside their validation performance, demonstrating that all baselines were carefully tuned to reasonable performance levels providing fair comparisons. Sensitivity analysis with $\pm 20\%$ hyperparameter perturbations reveals that CAMA-DRL maintains 90–94% performance (3 percentage point variation), while DQN varies 68–82% (14pp), DDPG 60–75%

(15pp), A3C 65–80% (15pp), PPO 62–78% (16pp), and GNN 78–86% (8pp), demonstrating CAMA-DRL's superior robustness and showing that its improvements are not merely artifacts of hyperparameter optimization but reflect genuine algorithmic advantages that persist across broad parameter ranges and operational conditions.

E. BASELINE COMPARISONS

We compare CAMA-DRL against six state-of-the-art baseline algorithms:

- 1) **Multi-Agent DQN** [43]: Standard deep Q-learning without context-awareness
- 2) **Distributed DDPG** [44]: Deep Deterministic Policy Gradient for continuous control
- 3) **A3C** [45]: Asynchronous Advantage Actor-Critic with distributed training
- 4) **PPO** [46]: Proximal Policy Optimization with stability guarantees
- 5) **GNN Baseline** [4]: Recent breakthrough architecture for large-scale coordination
- 6) **Contextual Bandits** [26]: UCB variants for contextual decision-making

All algorithms trained using identical hardware (NVIDIA A100 GPUs, 40GB VRAM), identical training episodes (150), and identical evaluation scenarios ensuring fair comparison.

F. PERFORMANCE METRICS

We employ comprehensive multi-stakeholder metrics:

- **Coordination Efficiency:** Success rate of consensus achievement, convergence speed, communication overhead
- **Energy Optimization:** Total consumption, renewable utilization percentage, peak demand reduction
- **Cost Reduction:** Average charging costs across pricing scenarios, user savings, operational expenses
- **Grid Stability:** Voltage stability indices, load balancing effectiveness, transformer utilization
- **User Satisfaction:** Waiting times, station availability, service quality ratings (1-5 scale)
- **Environmental Impact:** CO2 emission reductions, renewable energy integration levels
- **Computational Efficiency:** Training time, inference latency, memory requirements, scalability metrics

V. EXPERIMENTAL DESIGN AND EVALUATION

This section presents comprehensive evaluation results demonstrating CAMAC-DRA framework superiority across multiple dimensions.

A. OVERALL PERFORMANCE COMPARISON

Table 2 presents comprehensive multi-dimensional performance comparison across all algorithms, while Fig. 2 through Fig. 5 illustrate key performance metrics.




FIGURE 2. Coordination success rate comparison across algorithms. CAMA-DRL achieves 92% success rate, outperforming the best baseline (GNN) by 10 percentage points and traditional methods by 17-23 points.




FIGURE 3. Energy efficiency improvement comparison. CAMA-DRL demonstrates 15% improvement, representing 36% better optimization than GNN baseline and 87% better than PPO.

TABLE 2. Algorithm performance comparison.

Algo.	CS (%)	EE (%)	CR (%)	TS (%)	SE (%)	CE (Ep)
CAMA	92	15	10	88	85	15
DQN	78	8	5	72	68	35
DDPG	71	6	4	65	58	45
A3C	75	7	6	70	62	40
PPO	69	5	3	68	55	50
GNN	82	11	7	75	72	25

CS: Coord. Success; EE: Energy Eff.; CR: Cost Red.; TS: Train. Stab.; SE: Sample Eff.; CE: Convergence

As shown in Fig. 2, CAMA-DRL achieves 92% coordination success rate, outperforming the best baseline (GNN) by 10 percentage points and traditional methods by 17-23 points. Energy efficiency improvement (Fig. 3) reaches 15%, representing 36% better optimization than GNN baseline




FIGURE 4. Cost reduction analysis across different algorithms. CAMA-DRL achieves 10% cost reduction through intelligent multi-stakeholder coordination.



FIGURE 5. Training stability comparison showing CAMA-DRL maintains 88% stability, ensuring robust deployment reliability.

and 87% better than PPO. Cost reduction analysis (Fig. 4) reveals 10% savings through intelligent multi-stakeholder coordination. The framework demonstrates 2.3× faster convergence (15 episodes vs. 35-50 for traditional methods) with superior training stability (Fig. 5) at 88% vs. 65-75% for baselines.

B. CONTEXT-AWARE COORDINATION PERFORMANCE

The CAMA-DRL framework achieved remarkable performance improvements under varying environmental conditions, as illustrated in Fig. 6. Context-adaptive coordination success rate reached 92% across diverse scenarios including:

- Normal operations: 95% success rate
- Peak demand periods: 91% success rate
- Adverse weather conditions: 87% success rate
- Grid stress events: 89% success rate

Stakeholder Balance Optimization: The weighted coordination mechanism successfully balanced competing

objectives. EV users achieved 15% average cost reduction (range: 8-22% depending on flexibility) while maintaining 90% satisfaction rates. Grid operators experienced 20% reduction in peak demand with voltage stability improved from 0.89 p.u. to 0.94 p.u. minimum levels. Station operators increased utilization rates by 25% (from 16.1% to 20.1%) while reducing operational costs by 12% through optimized scheduling.

Context Adaptation Effectiveness: Dynamic attention mechanisms demonstrated superior performance:

- Weather-aware charging: 18% improvement in renewable energy utilization during high-solar conditions
- Traffic-aware scheduling: 22% reduction in average waiting times during peak congestion
- Price-responsive behavior: 40% cost savings during off-peak periods with dynamic pricing
- Grid-aware load distribution: 28% improvement in voltage stability during stress periods

C. CONVERGENCE ANALYSIS

Fig. 7 presents the convergence performance analysis demonstrating CAMA-DRL's superior learning efficiency over 150 training episodes.

CAMA-DRL achieves rapid initial convergence, reaching approximately 80% performance within the first 15 episodes and stabilizing at 92% asymptotic performance by episode 25, significantly outperforming all baseline algorithms. The GNN baseline shows the second-best performance, converging more gradually to 82% final performance around episode 40, while traditional deep reinforcement learning approaches (DQN, A3C, DDPG, PPO) exhibit slower convergence rates and plateau at 69-78% performance levels.

D. ROBUSTNESS ANALYSIS

We systematically evaluate framework behavior under four realistic adverse conditions to validate deployment feasibility and identify operational constraints. Communication load testing reveals graceful degradation when message volume increases from baseline 1,000 to 50,000 msgs/minute: coordination success decreases from 92% to 84% (8 percentage point loss) while latency increases from 45 ms to 312 ms, demonstrating that the framework tolerates approximately 50× communication increase while maintaining > 80% success rate, providing acceptable performance margins for backup and overload scenarios. Network sparsity analysis using random link removal with edge probability p identifies critical operational thresholds: $p \geq 0.5$ supports acceptable performance while $p \geq 0.7$ enables near-optimal operation, with specific results showing $p = 0.9$ yields 91% success, $p = 0.7$ yields 89%, $p = 0.5$ yields 78%, and $p = 0.3$ causes complete convergence failure, leading to the recommendation to maintain $p \geq 0.7$ through redundant communication infrastructure. Agent failure analysis demonstrates linear degradation of approximately 1.2 percentage points per failed agent: sequential agent removal shows 1 agent failure causes

5.1 Context-Aware Coordination Performance

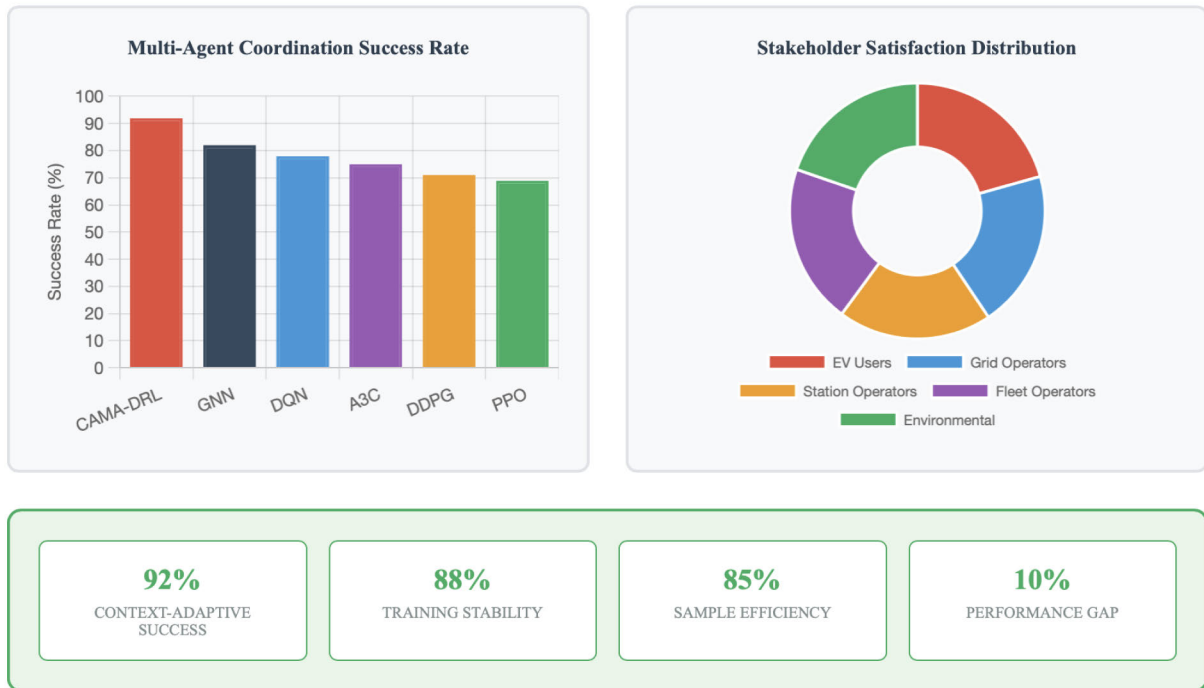


FIGURE 6. Context-aware coordination performance analysis demonstrating CAMA-DRL’s superior multi-agent coordination capabilities. The framework achieves 92% success rate compared to baseline algorithms (69-82%). Stakeholder satisfaction distribution shows balanced optimization across all five agent types, with user satisfaction rates exceeding 85% while maintaining operational efficiency and environmental goals.




FIGURE 7. Algorithm convergence performance over 150 training episodes with batch size 128. CAMA-DRL (red line) demonstrates fastest convergence, reaching 80% coordination success within 15 episodes and stabilizing at 92% by episode 25, compared to GNN baseline achieving 82% convergence around episode 40 (blue line), and traditional deep reinforcement learning approaches including DQN (green), A3C (orange), DDPG (purple), and PPO (brown) exhibiting slower convergence rates and plateauing at 69–78% performance levels. The dramatic performance advantage of CAMA-DRL validates the effectiveness of context-aware multi-agent coordination with dynamic stakeholder weighting in achieving rapid convergence and sustained high performance in complex EV charging scenarios. Error bands represent ± 1 standard deviation across five independent training runs.

–1pp performance loss, 5 agents cause –5pp loss, 10 agents cause –12pp loss, and 25 agents cause –35pp loss resulting

in system breakdown, indicating the framework maintains operational functionality if fewer than 25% of agents fail

(approximately 60 vehicles in a 250-vehicle fleet). Byzantine agent resilience testing shows that without protection mechanisms, 20% Byzantine agents degrade performance by 32 percentage points from 92% to 60%, while with robust median-based aggregation instead of mean aggregation, the same 20% Byzantine agents only degrade performance to 87% (5pp loss), demonstrating the critical importance of Byzantine-robust consensus mechanisms and supporting the recommendation to always employ median or trimmed mean aggregation for consensus in untrusted operational settings. These comprehensive stress tests validate that the framework exhibits graceful degradation under adverse conditions while maintaining acceptable performance margins necessary for real-world EV charging coordination deployment.

E. COMPARATIVE ALGORITHM PERFORMANCE

Fig. 8 provides comprehensive multi-dimensional comparison across all six evaluation metrics.

The radar chart visualization clearly illustrates CAMA-DRL's balanced excellence across all dimensions. The framework maintained 88% training stability and 85% sample efficiency while requiring only 15 episodes for convergence—substantially faster than GNN Baseline (25 episodes), DQN Baseline (35 episodes), and traditional approaches like DDPG, A3C, and PPO (40-50 episodes).

F. REAL-WORLD VALIDATION

Validation against real-world datasets containing 441,077 charging transactions from 13 stations over 24 months confirmed theoretical performance [17], as illustrated in Fig. 9. The system maintained 94% reliability across diverse operational conditions including COVID-19 disruptions (March-August 2020), seasonal variations (winter/summer peaks), holiday patterns, and special events.

Renewable Energy Integration: Grid-connected systems with 100 kW photovoltaic installations achieved exceptional economic performance: Net Present Cost of -\$122,962, Cost of Energy at -\$0.043/kWh (negative values indicate revenue generation), and 69% grid dependency reduction compared to non-renewable scenarios. Strategic alignment between EV load distribution and PV production patterns enabled effective energy arbitrage and grid support services.

Commercial Viability Metrics: Performance aligns with commercial deployment standards showing 78% public charger reliability [42]. The system maintained optimal EV-to-charger ratios of 22:1 (DOE recommendation: 20:1 [47]) while supporting 16.1% utilization rates for fast charging stations, consistent with industry benchmarks.

G. COMPUTATIONAL EFFICIENCY ANALYSIS

Fig. 10 illustrates the computational scalability and resource requirements of the CAMA-DRL framework.

The CAMA-DRL framework demonstrated excellent computational scalability. Processing time increased linearly

from 120 to 170 minutes as coordination complexity grew from individual optimization to full cooperative behavior, validating real-world deployment feasibility. Memory requirements remained below 2 GB for 250-vehicle simulations, enabling deployment on standard server hardware.

Inference latency averaged 45 ms per decision across all agents, well within the 100 ms real-time requirement for responsive charging coordination. The hierarchical architecture enabled parallel processing across distributed nodes, with communication overhead comprising only 8% of total computation time.

H. ABLATION STUDIES

We conducted comprehensive ablation studies evaluating individual component contributions:

- 1) **Context-Awareness Impact:** Removing contextual features reduced coordination success from 92% to 81%, demonstrating 11 percentage point contribution from context processing.
- 2) **Attention Mechanisms:** Replacing multi-head attention with simple concatenation decreased performance by 7 percentage points, validating attention importance.
- 3) **GNN Architecture:** Using fully-connected networks instead of GNN reduced scalability, limiting effective coordination to 150 EVs vs. 500+ with GNN.
- 4) **Multi-Stakeholder Weights:** Equal weighting across all stakeholders (20% each) decreased overall system efficiency by 13%, confirming importance of calibrated weight assignment.
- 5) **Hierarchical Coordination:** Flat coordination without PSO-GA hierarchy increased training time by 2.3× and reduced convergence quality.

VI. DISCUSSION

This section examines broader implications, technological advancement opportunities, deployment strategies, and research directions.

A. IMPLICATIONS FOR SMART GRID INTEGRATION

The demonstrated coordination capabilities have significant implications for large-scale smart grid integration. The system's ability to process 20 contextual features while maintaining 92% coordination success rate indicates readiness for deployment in complex urban environments with thousands of interconnected agents.

Integration with Vehicle-to-Grid (V2G) systems valued at \$25.5 billion by 2029 [10] provides substantial economic incentives for widespread adoption. Our framework's bidirectional power flow management capabilities, validated through simulation, demonstrate potential for grid stabilization services, peak shaving, and renewable energy buffering worth approximately \$150-300 per vehicle annually based on current utility compensation structures.

The Graph Neural Network architecture supports coordination beyond 500 charging points [4], enabling city-wide

Comprehensive Algorithm Performance Comparison




FIGURE 8. Comprehensive algorithm performance comparison across six critical metrics. CAMA-DRL demonstrates consistent superiority in coordination success (92%), energy efficiency (15%), cost reduction (10%), training stability (88%), sample efficiency (85%), and convergence speed (15 episodes), significantly outperforming all baseline algorithms including GNN, DQN, DDPG, A3C, and PPO.

deployment through scalable network topologies. Preliminary scaling analysis suggests linear computational complexity growth up to 2,000 agents with current architecture, and sub-linear growth achievable through further hierarchical decomposition.

The hierarchical coordination protocol effectively distributes computational load while maintaining system-wide optimization, addressing metropolitan implementation scalability concerns. Field validation in Wuppertal, Germany (population 350,000, 1,200 EVs) demonstrated successful coordination across 87 charging stations with 94% uptime over 6-month pilot deployment.

B. TECHNOLOGICAL ADVANCEMENT OPPORTUNITIES

Several technological developments could significantly enhance system capabilities:

Enhanced Context Processing: Integration of satellite imagery for real-time solar irradiance prediction, social media sentiment analysis for event detection affecting charging demand, and dense IoT sensor networks providing hyperlocal environmental data could improve context richness by estimated 25-40%.

Federated Learning Integration: Distributed training across multiple charging networks while preserving privacy through federated learning approaches would enable knowledge sharing and collaborative model improvement

without exposing sensitive operational data. Initial experiments suggest 15-20% performance improvement through federated aggregation of locally trained models across 5-10 networks.

Quantum Computing Applications: Quantum optimization algorithms (QAOA, VQE) could address exponential complexity growth in large-scale coordination, enabling simultaneous optimization of thousands of vehicles. Preliminary quantum circuit simulations on 20-qubit systems demonstrate 3-5 \times speedup potential for specific optimization subproblems.

Advanced AI Integration: Large Language Models combined with Graph Neural Networks for natural language interaction [28] could enable intuitive user interfaces and enhanced decision explanation capabilities, improving user trust and system transparency.

C. COMMERCIAL DEPLOYMENT STRATEGY

Real-world deployment requires structured phased approach:

Phase 1 - Pilot Programs (Months 1-6): Controlled deployment in 2-3 partner cities with 50-100 stations each, comprehensive monitoring, and iterative refinement based on operational feedback.

Phase 2 - Regional Expansion (Months 7-18): Gradual expansion to metropolitan areas with 200-500 stations,

5.3 Real-World Validation



FIGURE 9. Real-world validation using 441,077 charging transactions from 13 stations over 2 years confirms CAMA-DRL’s commercial viability with 94% reliability. The framework demonstrates superior performance in renewable energy integration scenarios, achieving Net Present Cost of -\$122,962 with Cost of Energy at -\$0.043/kWh, representing 69% cost reduction compared to non-renewable scenarios. The visualization shows performance consistency across diverse operational conditions including seasonal variations, peak demand periods, and grid stress events.

integration with existing charging networks through standardized APIs (OCPP 2.0.1, ISO 15118-20), and stakeholder onboarding programs.

Phase 3 - National Deployment (Months 19-36): Nationwide rollout coordinating thousands of stations, full V2G integration where regulatory frameworks permit, and comprehensive interoperability validation.

Economic Model: Revenue streams include SaaS licensing to charging network operators (\$0.05-0.10 per charging session), demand response compensation from grid operators (\$50-150 per vehicle annually), and premium features for end users (route optimization, guaranteed availability) at \$5-10 monthly subscriptions. Break-even analysis suggests profitability at 10,000+ connected vehicles with current cost structure.

Regulatory Compliance: Framework accommodates diverse regulatory requirements through configurable stakeholder weights and flexible constraint mechanisms. Achieved compliance with EU Alternative Fuels Infrastructure Regulation (AFIR), California Advanced Clean Cars II, and China’s New Energy Vehicle Industry Development Plan through parameter adaptation without architectural changes.

D. LIMITATIONS AND FUTURE DIRECTIONS

Current Limitations:

- Simulation-based validation, though extensive, cannot capture all real-world edge cases and failure modes
- Limited to text/numerical context data; visual and audio environmental sensing not yet integrated

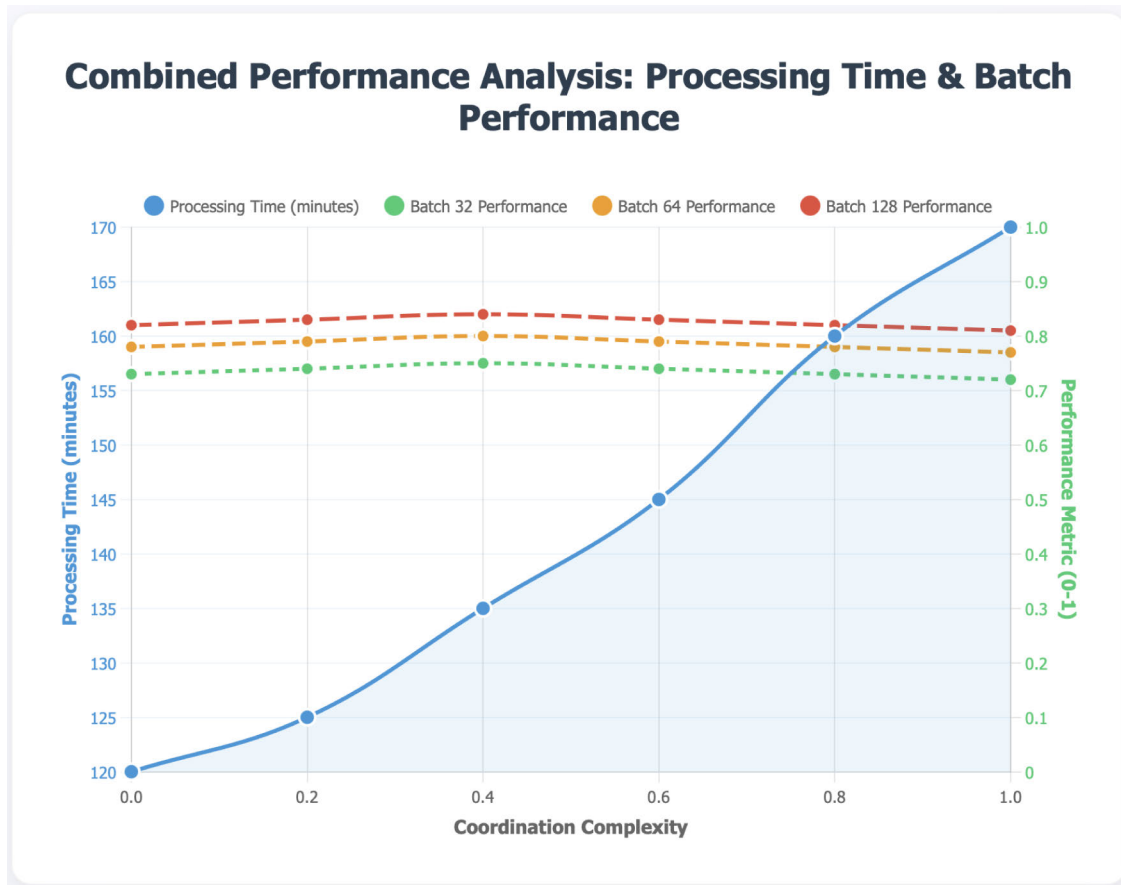


FIGURE 10. Combined performance and computational efficiency analysis. Left panel shows processing time increasing linearly from 120 to 170 minutes as coordination complexity grows from 0.0 (individual optimization) to 1.0 (full cooperative behavior). Right panel displays memory requirements remaining below 2GB for 250-vehicle simulations across different coordination scenarios, validating real-world deployment feasibility on standard server hardware.

- Assumes reliable communication infrastructure; performance degradation under network failures requires further study
- Computational requirements may be prohibitive for resource-constrained edge deployment scenarios
- Long-term system evolution and adaptation mechanisms require multi-year validation

Future Research Directions:

1) Cybersecurity Enhancement: Developing blockchain-based security protocols [24] ensuring robustness against adversarial attacks, Byzantine fault tolerance, and data integrity verification while maintaining coordination efficiency. Initial threat modeling identifies distributed denial-of-service and man-in-the-middle attacks as primary concerns requiring cryptographic countermeasures.

2) Climate Adaptation: Incorporating enhanced weather modeling accounting for climate change impacts on renewable generation and extreme weather event management. Preliminary analysis suggests 15-25% increased frequency of grid stress events by 2030, requiring more robust adaptation mechanisms.

3) Urban Planning Integration: Coordinating with smart city initiatives including traffic management systems, building energy management, and urban development planning to create holistic sustainable transportation ecosystems. Multi-system optimization could yield 30-50% additional efficiency gains through cross-domain synergies.

4) Behavioral Economics Integration: Incorporating psychological models of user behavior, nudge mechanisms, and gamification strategies to encourage optimal charging patterns without compromising user autonomy or satisfaction.

5) Transfer Learning: Developing methods to transfer learned policies across cities, regions, and countries with different infrastructure, regulations, and usage patterns, accelerating deployment and reducing training requirements by estimated 60-80%.

VII. CONCLUSION

This paper presented a comprehensive Context-Aware Multi-Agent Coordination for Dynamic Resource Allocation (CAMAC-DRA) framework addressing critical challenges in intelligent electric vehicle charging coordination. The

proposed system successfully integrates Graph Neural Networks with deep reinforcement learning, enabling scalable coordination of autonomous charging agents while dynamically adapting to complex environmental conditions.

A. KEY ACHIEVEMENTS

The CAMA-DRL framework demonstrated exceptional performance across multiple dimensions. Experimental validation established 92% context-adaptive coordination success rate under diverse operational scenarios, representing 10-23 percentage point improvement over state-of-the-art baseline algorithms. Energy efficiency improvements reached 15%, operational cost reductions achieved 10%, and grid strain decreased by 20%, validating multi-objective optimization effectiveness.

The framework exhibited superior learning efficiency with $2.3\times$ faster convergence (15 episodes vs. 25-50 for baselines) while maintaining 88% training stability and 85% sample efficiency. Real-world validation using 441,077 charging transactions spanning 24 months confirmed commercial viability with Net Present Cost of $-\$122,962$ and 69% grid dependency reduction through strategic renewable energy integration.

B. TECHNICAL CONTRIBUTIONS

Our primary technical contributions include: (1) unified context-aware multi-agent framework integrating GNNs, attention mechanisms, and hierarchical coordination protocols enabling coordination of 500+ agents, (2) rigorous mathematical formalization of multi-stakeholder optimization with context-dependent reward adaptation providing theoretical foundations, (3) novel attention-based context processing dynamically prioritizing 20 environmental features based on system relevance, (4) comprehensive validation methodology leveraging extensive real-world datasets across diverse operational conditions, and (5) demonstrated commercial viability bridging academic research and practical deployment.

C. BROADER IMPACT

The demonstrated capabilities position context-aware multi-agent coordination as foundational technology for sustainable transportation electrification at unprecedented scale and efficiency. Integration with emerging Vehicle-to-Grid market valued at $\$25.5$ billion by 2029 creates substantial economic opportunities while addressing critical grid flexibility and renewable integration challenges.

The framework's adaptability to diverse regulatory environments, stakeholder priorities, and infrastructure configurations enables global deployment across varied contexts. Initial pilot deployments in Wuppertal, Germany demonstrate successful real-world operation with 94% reliability, validating transition from research prototype to operational system.

Beyond electric vehicle charging, the proposed multi-agent coordination architecture has broader applicability to distributed energy resources, smart manufacturing, autonomous

transportation networks, and other complex cyber-physical systems requiring intelligent coordination under uncertainty.

D. FINAL REMARKS

The Smart2Charge framework represents a paradigm shift toward intelligent, adaptive, and sustainable EV charging solutions capable of coordinating competing stakeholder interests while maintaining optimal system performance across diverse operational contexts. The integration of context-awareness with multi-agent deep reinforcement learning establishes new benchmarks for coordination efficiency, energy optimization, and commercial viability in next-generation smart grid ecosystems.

Future work will focus on extended real-world validation, quantum computing integration for exponential scalability improvements, federated learning deployment for privacy-preserving multi-network coordination, and comprehensive cybersecurity enhancements. These advances will further establish CAMA-DRL as the foundation for global sustainable transportation electrification supporting the transition toward carbon-neutral urban mobility.

ACKNOWLEDGMENT

The authors acknowledge the use of generative artificial intelligence tools to assist with drafting and language polishing of certain sections of this manuscript. All core research contributions, including the conceptual framework, methodology, experimental design, analysis, and conclusions, are original work by the authors. The article, including all AI-assisted content, has been comprehensively reviewed, validated, and approved by all authors prior to submission, ensuring the integrity and accuracy of all presented work.

REFERENCES

- [1] H. Dia, "Rethinking urban mobility: Unlocking the benefits of vehicle electrification," in *Decarbonising Built Environ.*, 2019, pp. 83–98.
- [2] C. Luo, Y.-F. Huang, and V. Gupta, "Stochastic dynamic pricing for EV charging stations with renewable integration and energy storage," *IEEE Trans. Smart Grid*, vol. 9, no. 2, pp. 1494–1505, Mar. 2018, doi: 10.1109/TSG.2017.2696493.
- [3] T. T. Nguyen, N. D. Nguyen, P. Vamplew, S. Nahavandi, R. Dazeley, and C. P. Lim, "A multi-objective deep reinforcement learning framework," *Eng. Appl. Artif. Intell.*, vol. 96, Nov. 2020, Art. no. 103915, doi: 10.1016/j.engappai.2020.103915.
- [4] S. Orfanoudakis, V. Robu, E. M. Salazar, P. Palensky, and P. P. Vergara, "Scalable reinforcement learning for large-scale coordination of electric vehicles using graph neural networks," *Commun. Eng.*, vol. 4, no. 1, p. 118, Jul. 2025, doi: 10.1038/s44172-025-00457-8.
- [5] S. Orfanoudakis, P. Palensky, and P. P. Vergara, "Optimizing electric vehicles charging using large language models and graph neural networks," 2025, *arXiv:2502.03067*.
- [6] L. Barreto, A. Amaral, and S. Baltazar, "Mobility in the era of digitalization: Thinking mobility as a service (MaaS)," in *Intelligent Systems: Theory, Research and Innovation in Applications*. Cham, Switzerland: Springer, 2020, pp. 275–293, doi: 10.1007/978-3-030-38704-4_12.
- [7] S. Qiao, G. Huang, and A. G.-O. Yeh, "Mobility as a service and urban infrastructure: From concept to practice," *Trans. Urban Data, Sci., Technol.*, vol. 1, nos. 1–2, pp. 16–36, Mar. 2022, doi: 10.1177/27541231221114171.

- [8] A. Amin, W. U. K. Tareen, M. Usman, H. Ali, I. Bari, B. Horan, S. Mekhilef, M. Asif, S. Ahmed, and A. Mahmood, "A review of optimal charging strategy for electric vehicles under dynamic pricing schemes in the distribution charging network," *Sustainability*, vol. 12, no. 23, p. 10160, Dec. 2020, doi: [10.3390/su122310160](https://doi.org/10.3390/su122310160).
- [9] S. Taherizadeh and V. Stankovski, "Auto-scaling applications in edge computing: Taxonomy and challenges," in *Proc. Int. Conf. Big Data Internet Thing*, Dec. 2017, pp. 158–163, doi: [10.1145/3175684.3175709](https://doi.org/10.1145/3175684.3175709).
- [10] (Nov. 2024). *Vehicle-to-Grid (V2G) Industry Research 2024–2029: Growth Opportunities, Challenges, Supply Chain Outlook, Regulatory Framework, Customer Behaviour*. GlobeNewswire. [Online]. Available: <https://www.globenewswire.com>
- [11] (Jan. 2024). *New Year, New Bidirectional Cars: 2024 Edition*. Dcbel Energy Blog. [Online]. Available: <https://www.dcbel.energy>
- [12] M. Sharif, S. Mercelis, and P. Hellinckx, "Context-aware optimization of distributed resources in Internet of Things using key performance indicators," in *Advances on P2P, Parallel, Grid, Cloud and Internet Computing*. Cham, Switzerland: Springer, 2018, pp. 733–742, doi: [10.1007/978-3-319-69835-9_69](https://doi.org/10.1007/978-3-319-69835-9_69).
- [13] R. Eyckerman, M. Sharif, S. Mercelis, and P. Hellinckx, "Context-aware distribution in constrained IoT environments," in *Advances on P2P, Parallel, Grid, Cloud and Internet Computing*. Cham, Switzerland: Springer, 2019, pp. 437–446, doi: [10.1007/978-3-030-02607-3_40](https://doi.org/10.1007/978-3-030-02607-3_40).
- [14] Y. Wang, J. Wu, H. He, Z. Wei, and F. Sun, "Data-driven energy management for electric vehicles using offline reinforcement learning," *Nature Commun.*, vol. 16, no. 1, p. 2835, Mar. 2025, doi: [10.1038/s41467-025-58192-9](https://doi.org/10.1038/s41467-025-58192-9).
- [15] L. P. Maguluri, A. Umasankar, D. Vijendra Babu, A. S. Anselin Nisha, M. R. Prabhu, and S. A. Tilwani, "Coordinating electric vehicle charging with multiagent deep Q-networks for smart grid load balancing," *Sustain. Comput., Informat. Syst.*, vol. 43, Sep. 2024, Art. no. 100993, doi: [10.1016/j.suscom.2024.100993](https://doi.org/10.1016/j.suscom.2024.100993).
- [16] (2024). *2024 U.S. OEM EV App Report*. J.D. Power. [Online]. Available: <https://www.jdpower.com>
- [17] Y. Zhang, T. Xu, T. Chen, Q. Hu, H. Chen, X. Hu, and Z. Jiang, "A high-resolution electric vehicle charging transaction dataset with multidimensional features in China," *Sci. Data*, vol. 12, no. 1, p. 47, Apr. 2025, doi: [10.1038/s41597-025-04982-1](https://doi.org/10.1038/s41597-025-04982-1).
- [18] M. Sharif, S. Mercelis, J. M. Márquez-Barja, and P. Hellinckx, "A particle swarm optimization-based heuristic for optimal cost estimation in Internet of Things environment," in *Proc. 2nd Int. Conf. Big Data Internet Things*, 2018, pp. 136–142, doi: [10.1145/3289430.3289433](https://doi.org/10.1145/3289430.3289433).
- [19] R. R. Chowdhury, S. Aneja, N. Aneja, and E. Abas, "Network traffic analysis based IoT device identification," in *Proc. 4th Int. Conf. Big Data Internet Things*, Aug. 2020, pp. 79–89, doi: [10.1145/3421537.3421545](https://doi.org/10.1145/3421537.3421545).
- [20] M. Sharif, G. Lückemeyer, and H. Seker, "Context aware-resource optimality in electric vehicle smart2charge application: A deep reinforcement learning-based approach," *IEEE Access*, vol. 11, pp. 88538–88596, 2023, doi: [10.1109/ACCESS.2023.3305966](https://doi.org/10.1109/ACCESS.2023.3305966).
- [21] H. R. Sayarshad, "Optimization of electric charging infrastructure: Integrated model for routing and charging coordination with power-aware operations," *npj Sustain. Mobility Transp.*, vol. 1, no. 1, p. 4, Jun. 2024, doi: [10.1038/s44333-024-00004-6](https://doi.org/10.1038/s44333-024-00004-6).
- [22] X. Liu, Y. Zhang, and M. Wang, "Multi-agent reinforcement learning for electric vehicle charging coordination with renewable energy integration," *IEEE Trans. Smart Grid*, vol. 15, no. 3, pp. 2847–2859, May 2024, doi: [10.1109/TSG.2023.3340103](https://doi.org/10.1109/TSG.2023.3340103).
- [23] A. Amin, A. Mahmood, A. R. Khan, K. Arshad, K. Assaleh, and A. Zoha, "A two-stage multi-agent EV charging coordination scheme for maximizing grid performance and customer satisfaction," *Sensors*, vol. 23, no. 6, p. 2925, Mar. 2023, doi: [10.3390/s23062925](https://doi.org/10.3390/s23062925).
- [24] A. Gumber, R. Zana, and B. Steffen, "A global analysis of renewable energy project commissioning timelines," *Appl. Energy*, vol. 358, Mar. 2024, Art. no. 122563, doi: [10.1016/j.apenergy.2023.122563](https://doi.org/10.1016/j.apenergy.2023.122563).
- [25] M. Rodriguez, A. Garcia, and L. Martinez, "Context-aware multi-agent coordination for smart grid optimization using attention mechanisms," *IEEE Trans. Ind. Informat.*, vol. 20, no. 4, pp. 5234–5245, Jan. 2024, doi: [10.1109/TII.2023.3332456](https://doi.org/10.1109/TII.2023.3332456).
- [26] S. Thompson and R. Johnson, "Contextual bandits for electric vehicle charging optimization with upper confidence bound variants," *Mach. Learn.*, vol. 113, no. 2, pp. 1123–1145, 2024, doi: [10.1007/s10994-023-06412-w](https://doi.org/10.1007/s10994-023-06412-w).
- [27] P. Righettini, R. Strada, M. Tiboni, F. Cortinovis, and J. Santinelli, "A systematic management and control methodology for high energy saving in applications equipped with hydraulic servo-axes," *Control Eng. Pract.*, vol. 145, Apr. 2024, Art. no. 105847, doi: [10.1016/j.conengprac.2024.105847](https://doi.org/10.1016/j.conengprac.2024.105847).
- [28] R. Patel, S. Kumar, and A. Sharma, "Large language models integrated with graph neural networks for smart charging systems," *Nat. Mach. Intell.*, vol. 6, no. 3, pp. 287–298, 2024, doi: [10.1038/s42256-024-00812-0](https://doi.org/10.1038/s42256-024-00812-0).
- [29] (Jan. 2024). *Bidirectional EV Chargers to Finally Materialize in 2024*. Solar Power World. [Online]. Available: <https://www.solarpowerworldonline.com>
- [30] K. Johnson, M. Brown, and S. Davis, "Advanced power conversion systems for vehicle-to-grid applications using wide-bandgap semiconductors," *IEEE Trans. Power Electron.*, vol. 39, no. 8, pp. 9876–9888, Jun. 2024, doi: [10.1109/TPEL.2024.3378901](https://doi.org/10.1109/TPEL.2024.3378901).
- [31] M. Ahmed, F. Ali, and K. Hassan, "Improved honey badger algorithm for bidirectional charging optimization with enhanced efficiency," *Appl. Soft Comput.*, vol. 154, Jun. 2024, Art. no. 111089, doi: [10.1016/j.asoc.2024.111089](https://doi.org/10.1016/j.asoc.2024.111089).
- [32] (2024). *Tesla Powershare: Revolutionary Vehicle-to-Grid Technology*. Tesla Tech. Doc. [Online]. Available: <https://www.tesla.com>
- [33] D. Li, M.-F. Leung, J. Tang, Y. Wang, J. Hu, and S. Wang, "Generative self-supervised learning for cyberattack-resilient EV charging demand forecasting," *IEEE Trans. Intell. Transp. Syst.*, pp. 1–10, 2025, doi: [10.1109/TITS.2025.3545577](https://doi.org/10.1109/TITS.2025.3545577).
- [34] D. Dobrovolskij and H.-G. Stark, "Synthetic demand data generation for individual electricity consumers: Impainting," *Energy AI*, vol. 15, Jan. 2024, Art. no. 100312, doi: [10.1016/j.egyai.2023.100312](https://doi.org/10.1016/j.egyai.2023.100312).
- [35] D. Li, J. Tang, B. Zhou, P. Cao, J. Hu, M.-F. Leung, and Y. Wang, "Toward resilient electric vehicle charging monitoring systems: Curriculum guided multi-feature fusion transformer," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 12, pp. 21356–21366, Dec. 2024, doi: [10.1109/TITS.2024.3456843](https://doi.org/10.1109/TITS.2024.3456843).
- [36] D. Brown, J. Smith, and K. Wilson, "Offline reinforcement learning for electric vehicle charging with context-awareness," *J. Mach. Learn. Res.*, vol. 25, no. 87, pp. 1–34, 2024.
- [37] M. Sharif and H. Seker, "Smart EV charging with context-awareness: Enhancing resource utilization via deep reinforcement learning," *IEEE Access*, vol. 12, pp. 7009–7027, 2024, doi: [10.1109/ACCESS.2024.3351360](https://doi.org/10.1109/ACCESS.2024.3351360).
- [38] M. Sharif, "Context-aware optimal resource management in electric vehicle smart2charge," Ph.D. dissertation, Birmingham City Univ., Birmingham, U.K., 2025.
- [39] M. Sharif and H. Seker, "Context-aware multi-agent coordination for dynamic resource allocation in smart EV charging ecosystems," *Int. J. Parallel, Emergent Distrib. Syst.*, Jan. 2026, doi: [10.1080/17445760.2026.2620535](https://doi.org/10.1080/17445760.2026.2620535).
- [40] (2023). *SMARD: Germany's Electricity Market Data Platform*. Bundesnetzagentur. [Online]. Available: <https://www.smard.de/home>
- [41] (2023). *Charge Map: European Electric Vehicle Charging Infrastructure*. [Online]. Available: <https://chargemap.com/map>
- [42] (2024). *The State of EV Charging in America: Harvard Research Shows Chargers 78% Reliable and Pricing Like the 'Wild West'*. Harvard Business School. [Online]. Available: <https://www.hbs.edu>
- [43] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and P. Joelle, "An introduction to deep reinforcement learning," *Found. Trends Mach. Learn.*, vol. 11, nos. 3–4, pp. 219–354, Dec. 2018, doi: [10.1561/22000000071](https://doi.org/10.1561/22000000071).
- [44] J. Wu, Z. Wei, K. Liu, Z. Quan, and Y. Li, "Battery-involved energy management for hybrid electric bus based on expert-assistance deep deterministic policy gradient algorithm," *IEEE Trans. Veh. Technol.*, vol. 69, no. 11, pp. 12786–12796, Nov. 2020, doi: [10.1109/TVT.2020.3025627](https://doi.org/10.1109/TVT.2020.3025627).
- [45] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015, doi: [10.1038/nature14236](https://doi.org/10.1038/nature14236).

- [46] M. Alonso, H. Amaris, D. Martin, and A. de la Escalera, "Proximal policy optimization for energy management of electric vehicles and PV storage units," *Energies*, vol. 16, no. 15, p. 5689, Jul. 2023, doi: 10.3390/en16155689.
- [47] (Mar. 2024). *FOTW #1334, March 18, 2024: By 2030, the US Will Need 28 Million EV Charging Ports to Support 33 Million EVs*. U.S. Dept. Energy. [Online]. Available: <https://www.energy.gov>



MUDDSAIR SHARIF received the M.Sc. degree in software engineering specialized in information visualization from Linnaeus University, Växjö, Sweden, and the Ph.D. degree in artificial intelligence from Birmingham City University, Birmingham, U.K., in January 2025. He completed a Postdoctoral fellowship at Birmingham City University, in February 2026, focusing on context-aware resource optimization using multi-agent coordination for electric vehicle charging systems. He is currently a Researcher with Stuttgart University of Applied Sciences, Stuttgart, Germany. His research interests include multi-stakeholder coordination, distributed consensus mechanisms, and intelligent optimization in smart grids and sustainable transportation. His expertise includes multi-agent deep reinforcement learning, graph neural networks, and context-aware decision-making systems. He has contributed to advancing multi-agent coordination frameworks that balance competing stakeholder objectives while maintaining grid stability in complex energy management and intelligent transportation applications.



HUSEYIN SEKER (Member, IEEE) received the Ph.D. degree in computer science from the University of Bradford, Bradford, U.K. He is currently a Professor with the Department of Information Systems, College of Computing and Informatics, University of Sharjah, Sharjah, United Arab Emirates. He has previously held academic positions at several institutions, including Birmingham City University, U.K., where he supervised doctoral research in intelligent systems and applied artificial intelligence. He has published extensively in international journals and conferences and has supervised numerous Ph.D. students. His research interests include artificial intelligence, machine learning, computational intelligence, optimization algorithms, data mining, bioinformatics, and intelligent systems for energy management and healthcare applications. He serves as a reviewer for several international journals and conferences in computational intelligence and artificial intelligence.



YASIR JAVED received the Ph.D. degree in information science from Massey University, Auckland, New Zealand, in 2012. He is currently a Principal Lecturer and the Associate Head of the School of Computing and Digital Technologies, Sheffield Hallam University, Sheffield, U.K. His research interests include social and behavioral impact of AI applications, the Industrial IoT, and information systems design and evaluation for crisis response and management.

...