Multi-agent AI Systems for Autonomous Smart Grid Operation and Renewable Energy Coordination

Kwadwo Larbi
Department of Renewable Engineering, University of Mines and Technology, Tarkwa, Ghana

Email: larbikwadwo123@gmail.com
Manuscript received August 6, 2026; accepted September 16, 2026; published November 20, 2026.
doi: 10.12720/sgce.15.2.123-145

Abstract: This systematic review synthesizes evidence on multi-agent artificial intelligence for autonomous smart-grid operation and renewable-energy coordination, covering publications from 1 January 2023 to 31 August 2026. Following PRISMA 2020 and PRISMA-S, 39 studies were included in the qualitative synthesis. The review distinguishes traditional Multi-agent Systems (MAS), Multi-agent Reinforcement Learning (MARL), generative AI-augmented MAS, and emerging Agentic AI. Peer-reviewed evidence supports applications in voltage and frequency regulation, distributed energy management, demand response, electric-vehicle coordination, and service restoration, although most evaluations remain simulation-based and heterogeneous. Preprints are explicitly identified and are used mainly for rapidly evolving Agentic AI, LLM-agent, communication, and synthetic-data topics. Excluding preprints does not materially change the conclusions for established MAS/MARL applications, but it substantially weakens claims of deployment-ready Agentic AI. The most credible deployment pathway is therefore a layered architecture combining deterministic safety constraints, distributed learning, validated engineering solvers, cybersecurity controls, and human oversight. Regulatory compliance, market rules, lifecycle economics, workforce readiness, data governance, standardized benchmarks, and field validation remain decisive barriers to scalable adoption.

Keywords: multi-agent systems, agentic AI, smart grids, multi-agent reinforcement learning, renewable energy coordination

1. Introduction

The global electricity sector is undergoing structural change driven by decarbonization, decentralization, digitalization, and electrification. Solar photovoltaic and wind generation are expanding rapidly, reshaping power-system planning and operation [1, 2]. Grid expansion, control infrastructure, and flexibility, however, have not advanced at the same pace; electricity networks can therefore become a bottleneck to secure energy transitions [3]. High renewable penetration creates operational challenges that extend beyond conventional balancing, including rapid net-load changes, forecast errors, congestion, voltage deviations, frequency-control difficulties, reverse power flows, and renewable curtailment. These challenges are intensified by the growth of batteries, electric vehicles, flexible buildings, prosumers, energy communities, aggregators, and interconnected microgrids. Multi-energy management problems are increasingly multidimensional, multi-objective, and partially observable [4], while coordination across energy carriers and distributed resources requires decision mechanisms that can capture nonlinear interactions and respond to changing system states [5].

Conventional supervisory control, rule-based scheduling, model predictive control, and optimal power

flow remain indispensable, but they may depend on accurate mathematical models, broad system observability, reliable communication, and repeated computationally intensive optimization. Centralized artificial-intelligence methods can improve prediction and control but may introduce communication bottlenecks, privacy concerns, limited scalability, and single points of failure [6]. Multi-agent artificial intelligence offers a complementary architecture in which computational agents represent generators, storage systems, flexible loads, electric vehicles, microgrids, network operators, or market participants. Each agent observes relevant conditions, selects actions within a defined authority boundary, and exchanges information to pursue local and system-wide objectives. Peer-reviewed studies have applied multi-agent learning to demand response in grid-responsive buildings [7], autonomous renewable microgrids [8], residential hybrid energy systems [9], and multi-objective microgrid dispatch [10]. Under centralized training and decentralized execution, agents can use shared information during training while acting locally during operation. Collectively, these results support the use of MAS and MARL for renewable coordination, storage scheduling, demand response, voltage regulation, electric-vehicle charging, peer-to-peer trading, congestion management, service restoration, and resilient islanded operation.

Recent peer-reviewed studies also extend multi-agent coordination from isolated microgrids to communities and sector-coupled environments. Community-level MARL has been used to coordinate vehicle-to-grid resources, shared storage, household demand, and local renewable generation [11]; multi-agent demand response has been applied to energy-intensive manufacturing [12]; and a unified Markov-game formulation has integrated photovoltaic design, battery degradation, comfort preservation, and peer-to-peer exchange [13]. Important limitations remain, including non-stationarity, partial observability, unstable training, reward misspecification, unsafe exploration, communication delay, cyberattacks, privacy leakage, explainability, transferability, and compliance with network constraints. Accordingly, this review has four objectives: to classify multi-agent AI architectures, learning paradigms, coordination mechanisms, communication structures, and agent roles; to synthesize applications in renewable forecasting, energy management, storage control, demand response, voltage and frequency regulation, electric-vehicle coordination, market participation, resilience, and restoration; to evaluate the evidence against technical, economic, environmental, computational, scalability, privacy, cybersecurity, explainability, and safety criteria; and to identify research priorities for trustworthy, interoperable, and deployment-ready grid intelligence. Because evidence maturity differs sharply across paradigms, peer-reviewed studies, preprints, and standards or agency reports are distinguished throughout the review.

The review makes four principal contributions:

i. It develops an integrated taxonomy linking agent types, grid layers, control timescales, information-sharing structures, learning algorithms, and operational objectives.
ii. It comparatively evaluates centralized, decentralized, hierarchical, federated, cooperative, competitive, and hybrid multi-agent architectures, clarifying their assumptions, strengths, and failure modes.
iii. It connects algorithmic performance with power-system feasibility by emphasizing physical constraints, stability, uncertainty, resilience, communication reliability, cybersecurity, explainability, market and regulatory constraints, and human oversight.
iv. It proposes a forward-looking roadmap centered on physics-informed learning, digital twins, safe reinforcement learning, uncertainty quantification, secure coordination, standardized benchmarks, hardware-in-the-loop testing, data governance, and staged field deployment.

The remainder of the paper presents the review methodology, defines the principal multi-agent paradigms, develops a comparative taxonomy, synthesizes applications and enabling technologies, evaluates technical and socio-economic deployment barriers, and identifies priorities for trustworthy autonomous operation of renewable-rich power systems.

2. Methodology

2.1. Review Design

This study adopted a structured systematic-review approach to identify, classify, and synthesize recent research on multi-agent artificial intelligence for autonomous smart-grid operation and renewable-energy coordination. Study identification, screening, eligibility assessment, and reporting were organized according to PRISMA 2020 [14]. Search reporting followed PRISMA-S recommendations [15], with the search concepts, eligibility window, source-specific syntax, and supplementary-search procedures documented to support reproducibility.

2.2. Information Sources and Search Strategy

The eligible publication window was 1 January 2023 through 31 August 2026. The formal database search covered Scopus, Web of Science Core Collection, IEEE Xplore, and ScienceDirect. SpringerLink and the ACM Digital Library were used as supplementary publisher-platform checks to verify coverage and retrieve publisher versions; overlapping records located through those platforms were not counted as separate identification streams. Google Scholar and backward reference searching were used to identify additional records. The search was constructed from three concept families: (1) multi-agent intelligence, (2) power-system context, and (3) operational function. The fully parenthesized expression below is the cross-platform conceptual template. Appendix A reports the executable translation for each source, including platform-specific adaptations where interface limits prevented the complete template from being entered as a single query.

("multi-agent system" OR "multi-agent systems" OR "multiagent system" OR "multiagent systems" OR "multi-agent reinforcement learning" OR MARL OR "Agentic AI" OR "LLM agent" OR "LLM agents" OR "large language model agent" OR "large language model agents" OR ("generative AI" AND agent)) AND ("smart grid" OR "smart grids" OR microgrid OR microgrids OR "power system" OR "power systems" OR "distribution network" OR "distribution networks" OR "distributed energy resource" OR "distributed energy resources" OR DER OR "renewable energy") AND (control OR coordination OR dispatch OR scheduling OR "energy management" OR "demand response" OR restoration OR "voltage control" OR "frequency control" OR trading).

2.3. Eligibility Criteria

Studies were eligible when they were published in English within the specified date window, explicitly addressed MAS, MARL, Agentic AI, LLM-based multi-agent systems, or generative AI-augmented agents, and focused on smart grids, microgrids, distributed energy resources, renewable-energy coordination, electricity markets, or related grid operations. Primary peer-reviewed journal articles and conference papers were eligible when they presented an architecture, algorithm, experiment, simulation, benchmark, or implementation framework with sufficient methodological detail for evidence extraction. Preprints were eligible only when they addressed rapidly emerging Agentic AI or closely related multi-agent capabilities not yet adequately represented in peer-reviewed grid literature; preprints were never used as the sole basis for claims of deployment readiness. Review articles, standards, and government or laboratory reports were retained as contextual sources for taxonomy, regulation, or field framing but were not counted as primary performance studies in the PRISMA total. Editorials, news items, duplicate reports, inaccessible full texts, studies without an explicit multi-agent component, and papers without a clear power-system or grid-operational connection were excluded.

2.4. Study Selection and Quality Assessment

The PRISMA flow is shown in Fig. 1. The search identified 2,136 records through the four core databases and 25 additional records through Google Scholar and backward citation searching. After 1,976 duplicate or

near-duplicate index records were removed, 185 unique records entered title-and-abstract screening. Of these, 121 were excluded for failing the topical or study-design criteria, leaving 64 reports for full-text eligibility assessment. Twenty-five full-text reports were excluded, and 39 studies were retained for the qualitative synthesis. Duplicate detection used DOI and exact-title matching followed by manual inspection of near-duplicate titles, author lists, and publication years. Screening proceeded in two stages: a deliberately inclusive title-and-abstract screen followed by full-text assessment against all eligibility criteria. The review was conducted by a single author; therefore, inter-reviewer agreement statistics were not applicable, and the absence of independent duplicate screening is recognized as a potential source of selection bias.

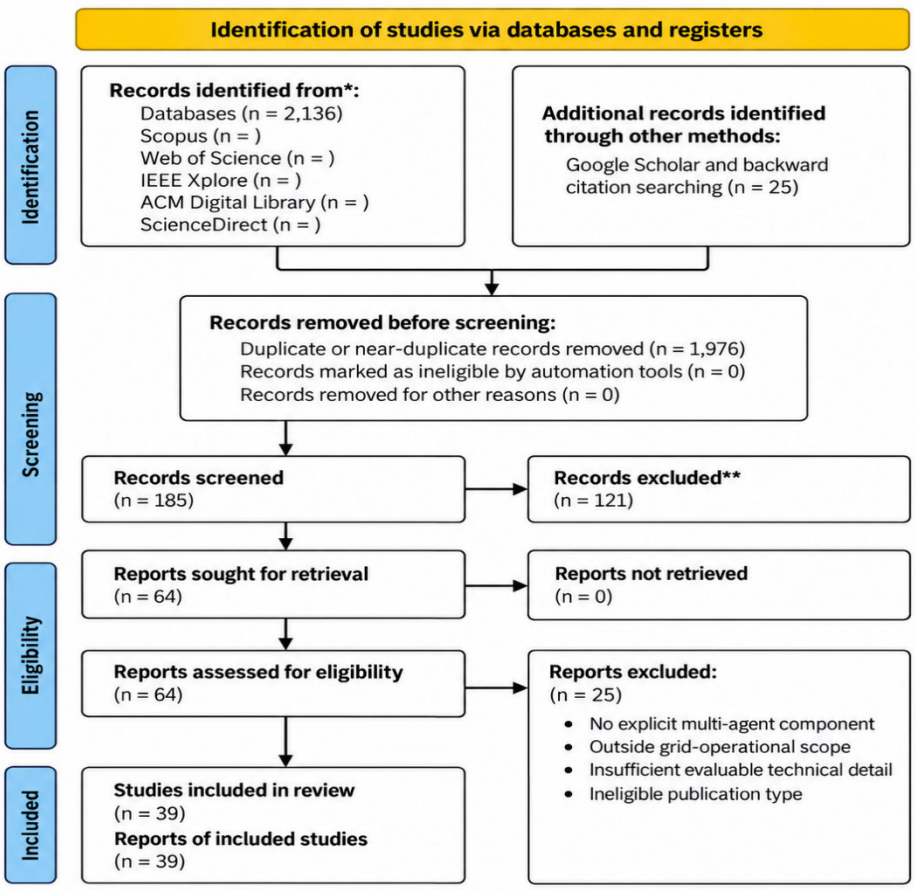


Fig. 1. PRISMA 2020 flow diagram for study identification, screening, eligibility assessment, and inclusion.

Methodological quality was assessed with an eight-domain appraisal tailored to computational power-system studies: (1) clarity and relevance of the research question; (2) reproducibility of the agent architecture, learning algorithm, state/action definitions, and coordination mechanism; (3) transparency of datasets, scenarios, and test systems; (4) adequacy of experimental design and comparator baselines; (5) appropriateness and completeness of performance metrics and, where applicable, statistical reporting; (6) representation of physical grid constraints, stability requirements, or safety limits; (7) robustness testing under uncertainty, topology change, communication impairment, cyber disturbance, or distribution shift; and (8) discussion of reproducibility, limitations, external validity, and deployment assumptions. Each domain was rated 2 (adequately addressed), 1 (partially addressed or incompletely reported), or 0 (not addressed or not reproducible); domains that were genuinely inapplicable were marked N/A and removed from the denominator. Normalized scores of at least 80%, 60–79%, and below 60% were interpreted as higher-, moderate-, and lower-confidence evidence, respectively.

Quality ratings were used to weight the narrative synthesis rather than as a rigid exclusion threshold because the included studies span control experiments, simulations, benchmarks, conference implementations, and architecture papers that are not directly commensurable. Lower-confidence studies could inform research gaps or hypothesis generation but were not treated as independent evidence of safety, scalability, economic viability, or field readiness. Publication status was recorded separately from methodological quality so that peer review was not conflated with study design; nevertheless, unreviewed preprints were assigned lower evidentiary certainty when conclusions depended on claims not corroborated by peer-reviewed sources.

2.5. Data Extraction and Evidence Synthesis

Data were extracted with a structured framework covering publication year, publication status, MAS paradigm, agent roles, architectural structure, learning algorithm, communication mechanism, grid application, renewable resources, test system, evaluation metrics, principal contributions, reported limitations, and quality-appraisal judgments. Evidence was grouped into foundational concepts, architectural paradigms, application domains, enabling technologies, operational challenges, socio-economic and governance barriers, and future research directions. Because the studies differed substantially in algorithms, networks, datasets, objectives, and evaluation measures, statistical meta-analysis was not appropriate. A comparative thematic synthesis was therefore used to identify methodological patterns, performance trends, evidence gaps, and requirements for trustworthy deployment.

2.6. Publication Status and Sensitivity Analysis

Every source used as empirical or technical evidence was coded as peer reviewed, preprint, or contextual grey literature/standard. All preprint citations are explicitly identified as such in the text and tables. The primary synthesis privileges peer-reviewed evidence. A qualitative sensitivity analysis was then conducted by removing all preprints and reassessing each thematic conclusion as unchanged, weakened, or dependent on preprint evidence. Conclusions that remained supported by peer-reviewed journal or conference studies were retained without qualification; conclusions materially weakened by preprint removal were narrowed and presented as emerging evidence or future research rather than established capability.

3. Foundational Concepts and Multi-Agent Architectures

3.1. Terminology and Relationship Taxonomy

Artificial-Intelligence (AI) agent: A computational entity that receives observations or information, maintains task-relevant state, selects actions toward an objective, and operates within an explicitly defined authority boundary. An agent may be rule based, optimization based, learning based, generative-model based, or hybrid.

Multi-Agent System (MAS): A system containing two or more interacting agents whose local actions can affect one another or a shared environment. MAS is the broad architectural category in this review; it does not imply reinforcement learning or the use of a large language model.

Multi-Agent Reinforcement Learning (MARL): A learning-based subset of MAS in which multiple agents improve policies through reward-guided interaction. MARL may be cooperative, competitive, or mixed and may use centralized training with decentralized execution. MARL therefore describes a learning mechanism within a multi-agent architecture, not a synonym for MAS.

Generative AI-augmented MAS: A conventional MAS or MARL architecture in which a generative model, such as a large language model or generative multimodal model, supports one or more agent functions through retrieval, summarization, scenario generation, code generation, interface translation, or decision support. The presence of a generative model does not by itself make the system agentic.

LLM-based multi-agent system: A multi-agent configuration in which one or more agents use a large language model as a principal reasoning or communication component. Such a system may be generative AI-augmented, agentic, both, or neither, depending on whether it has persistent goals, tool-mediated action, memory, and closed-loop replanning. A peer-reviewed survey provides a general taxonomy of LLM-based multi-agent systems [16], while a broader survey used here for conceptual triangulation remains a preprint [17].

Agentic AI: In this review, Agentic AI denotes a goal-directed AI system that can decompose multi-step tasks, maintain contextual state or memory, invoke external tools or models, evaluate intermediate results, and revise plans with limited human intervention. Agentic AI can be single-agent or multi-agent; therefore, Agentic AI and MAS overlap but neither is a strict subset of the other. For power systems, the term is restricted here to bounded supervisory, analytical, planning, and workflow-orchestration functions unless a study provides explicit validated control authority.

Relationship among terms: MAS is the umbrella coordination architecture; MARL is one learning approach that can implement MAS; generative AI is a model capability that can augment MAS or MARL; LLM-based multi-agent systems are a particular generative-model implementation; and Agentic AI describes a level and pattern of goal-directed autonomy that may use generative models and may be implemented as one or many agents. This separation prevents the common but misleading practice of treating MARL, LLM agents, and Agentic AI as interchangeable concepts.

3.2. Core Principles of Multi-agent Systems

Within this taxonomy, effective MAS operation depends on autonomy, reactivity, proactiveness, and social ability. A general MAS survey emphasizes explicit perceptions, objectives, decision rules, and communication interfaces as the basis of operational agency [18]. Reactivity allows an agent to respond to disturbances such as renewable forecast errors, voltage deviations, and abrupt load changes, while proactiveness enables preventive actions and longer-horizon optimization. Social ability supports cooperation, negotiation, information exchange, and conflict resolution across heterogeneous devices and stakeholders. Peer-reviewed smart-grid studies show that these properties can support hierarchical monitoring and control [19] and can be combined with adaptive learning and collaborative decision making [20]. In renewable-rich grids, however, autonomy must remain bounded by physical constraints, cybersecurity policies, operator authority, and interoperability requirements [21, 22].

3.3. Evolution from Traditional MAS to MARL and Agentic AI

Traditional MAS relied on rule-based reasoning, expert systems, auctions, consensus protocols, game-theoretic mechanisms, and belief-desire-intention models. These approaches provide explicit roles and interpretable coordination but can become brittle when grid dynamics, participant behavior, or renewable uncertainty depart from predefined assumptions. MARL shifts the emphasis toward experience-driven coordination by allowing agents to learn policies from interactions and rewards; however, scalability, non-stationarity, credit assignment, robustness, and communication remain substantial challenges [23]. LLM-based multi-agent systems add semantic reasoning and natural-language collaboration [16], while planning, tool use, reflection, and memory are emphasized in two general-agent surveys that remain preprints [24, 25]. Grid-specific Agentic AI evidence is newer: peer-reviewed studies and reviews describe customized large-model agents and bounded autonomous frameworks [22, 26, 27], but the number of field-validated implementations is still small. Accordingly, this review treats Agentic AI as an emerging supervisory and orchestration paradigm rather than as a mature replacement for MARL, deterministic optimization, or protection-grade control.

3.4. Hierarchical, Distributed, and Hybrid Architectural Paradigms

Multi-agent smart-grid architectures can be organized as hierarchical, distributed, or hybrid systems. Hierarchical architectures assign authority across device agents, microgrid coordinators, aggregators, market operators, and distribution-system operators. This decomposition can align decisions across operational timescales but may create upper-layer bottlenecks or points of failure [19]. Distributed architectures allow peer agents to negotiate, optimize, or learn from local observations, improving modularity, privacy, and fault tolerance, although performance depends on reliable communication and compatible objectives [28]. Hybrid architectures combine local autonomy with centralized training, supervisory optimization, distributed model predictive control, or market clearing. Reviews and implementations of peer-to-peer energy management, MARL, and distributed MPC illustrate this layered approach [21, 29, 30]. Architectural selection should therefore reflect grid size, latency, observability, privacy, communication availability, failure tolerance, and the consequences of unsafe decisions rather than algorithmic novelty alone.

3.5. Enabling Technologies for Intelligent Multi-agent Coordination

Four enabling technologies are particularly relevant to multi-agent grid coordination. Digital twins provide updated virtual representations of assets and networks for state estimation, contingency rehearsal, policy testing, and predictive maintenance [31, 32]. Edge intelligence moves inference and selected learning functions closer to substations, inverters, meters, buildings, and flexible loads, reducing latency, bandwidth dependence, and raw-data exposure, although resource constraints and device management remain important [33]. Blockchain can provide tamper-evident records, identities, and settlement support for transactive energy, but consensus overhead is generally unsuitable for fast inner-loop control [34]. Federated learning allows agents, prosumers, or utilities to train shared models without centralizing all raw data; applications include dynamic pricing and demand response [35] and security-oriented distributed learning [36]. None of these technologies removes the need for explicit safety constraints, interoperability, lifecycle maintenance, and governance.

4. Comparative Taxonomy of Multi-agent System Paradigms

Multi-agent intelligence in smart grids has progressed from explicitly programmed coordination toward adaptive learning and, more recently, cognitively orchestrated decision support. For analytical clarity, Table 1 uses the definitions established in Section 3: traditional MAS denotes explicitly programmed coordination; MARL denotes reward-driven multi-agent policy learning; generative AI-augmented MAS denotes conventional MAS or MARL enhanced by a generative model; and Agentic AI denotes bounded goal-directed planning, tool use, memory, evaluation, and replanning. The categories can overlap, and the table therefore compares their dominant decision mechanisms rather than treating them as mutually exclusive technologies.

Table 1 indicates that no single paradigm provides sufficient speed, adaptability, interpretability, scalability, and safety across all grid-control timescales. Traditional MAS remains appropriate for deterministic and tightly bounded functions, whereas MARL is better aligned with repetitive distributed optimization under uncertainty. Generative AI-augmented MAS can improve knowledge access, scenario generation, and human-machine interaction. Agentic AI has potential for supervisory workflow orchestration, contingency analysis, and operator support, but the grid-specific evidence base is considerably less mature and should not be interpreted as evidence for unrestricted real-time control. Operationally credible systems are therefore more likely to combine fast local controllers, adaptive MARL policies, deterministic power-system solvers, and supervised generative or agentic interfaces.

Table 1. Comparative Taxonomy and Operational Mapping of Multi-agent System Paradigms

Comparison dimension	Traditional MAS	MARL	Agentic AI	Generative AI-augmented MAS	Mapping to grid operational requirements
Core decision mechanism	Rules, logic, contracts, auctions, consensus, or predefined optimization [37].	Policies learned through rewards and repeated environmental interaction [38].	Goal decomposition, reasoning, memory, tool execution, and feedback loops; grid-specific evidence is emerging [26].	Generative models support existing agent decisions through prompts, retrieval, simulation, or code generation [39].	Determines whether operation requires deterministic control, adaptive optimization, or cognitive assistance.
Agent autonomy	Bounded by predefined roles and protocols.	High within the learned state and action spaces.	Potentially high task autonomy, but appropriate grid use requires explicit supervisory and permission constraints.	Usually advisory unless connected to executable controllers.	Protection and emergency control require bounded autonomy; planning can tolerate greater autonomy.
Adaptation to change	Limited unless rules or models are manually updated.	Learns renewable, demand, price, and device dynamics from experience [9].	Replans when tools, context, goals, or operating conditions change.	Updates recommendations through retrieved information or newly generated scenarios.	MARL and agentic approaches better support uncertain renewable and demand conditions.
Coordination model	Contract-net, negotiation, voting, auction, consensus, or supervisory coordination.	Cooperative, competitive, mixed, centralized-training, or decentralized-execution structures.	Role-specialized agents may collaborate through planning and task delegation [16].	Generative models translate, summarize, or mediate information among heterogeneous agents.	Supports DER aggregation, market coordination, microgrid cooperation, and control-room integration.
Planning capability	Restricted to encoded plans, schedules, or optimization routines.	Primarily reward-driven; long-horizon planning depends on policy and reward design.	Explicit multi-step planning, execution monitoring, reflection, and replanning; practical grid validation remains limited.	Generates candidate schedules, scenarios, contingencies, or control formulations.	Potential use in restoration planning and cross-functional workflows; engineering validation is required before action execution.
Knowledge representation	Ontologies, databases, state machines, expert rules, and network models.	Neural state representations and learned value or policy functions.	Context memory, domain knowledge bases, retrieval systems, and tool descriptions.	Large language or multimodal models can connect structured and unstructured grid information [40].	Enables integration of SCADA measurements, manuals, market rules, alarms, and operator knowledge.

Real-time responsiveness	Fast when rules are simple and locally executable.	Fast during inference, although training is computationally expensive.	Potentially slower because of iterative reasoning and external tool calls; unsuitable for protection-grade timing without a separate verified controller.	Suitable mainly for supervisory or advisory timescales.	Traditional MAS and trained MARL are better aligned with fast control; LLM-based systems fit slower advisory and supervisory layers.
Physical constraint handling	Constraints are explicitly encoded and therefore enforceable.	Requires constrained learning, safety layers, shielding, or verified action projection [38].	Must delegate numerical validation to trusted engineering solvers.	Can formulate constraints but should not independently certify feasibility.	Voltage, frequency, thermal, reserve, and stability limits require deterministic verification.
Scalability	Modular, but communication and rule complexity increase with system size.	Scales through decentralized execution, although non-stationarity and joint action spaces remain problematic.	Functional specialization may support modular expansion [26], but serving cost and coordination overhead can increase with agent count.	Reduces human workload but may increase model-serving and communication costs.	Large DER populations require distributed execution and selective information exchange.
Explainability	Relatively high because decisions follow explicit rules.	Often limited because policies are encoded in neural networks.	Produces natural-language rationales, although explanations may not faithfully represent internal processing.	Strong for summarization, reporting, alarm interpretation, and operator communication.	High-stakes actions require traceable evidence, solver outputs, logs, and human review [41].
Data and computational demand	Low to moderate; depends mainly on model complexity.	High during training and moderate during online inference.	High because reasoning, retrieval, memory, and tool orchestration consume substantial resources.	High for model training or cloud inference but lower when compact domain models are used.	Edge control favours lightweight agents; control centres can support more computationally intensive systems.
Robustness and uncertainty	Reliable within anticipated scenarios but brittle under unmodelled events.	Learns stochastic behaviour but may fail under distribution shifts.	Can compare scenarios and revise plans, but may hallucinate or misuse tools.	Generates diverse scenarios and supplements incomplete information, but outputs require validation.	Hybrid architectures are preferable for renewable uncertainty, contingencies, and rare operating states.
Cybersecurity and privacy	Smaller learning attack surface, but agent communication remains	Exposed to data poisoning, adversarial observations, policy manipulation, and privacy leakage.	Adds prompt injection, tool misuse, memory corruption, and excessive-permission risks.	May expose sensitive operational data through external models or	Zero-trust access, local processing, authentication, least privilege, and audit trails

	vulnerable.			generated content.	are essential.
Principal grid applications	Protection coordination, restoration rules, market negotiation, distributed scheduling, and microgrid control [37].	Voltage control, storage dispatch, demand response, frequency support, trading, and renewable coordination [10, 42, 43, 44].	Early peer-reviewed evidence supports analytical and orchestration functions [26, 27]; a power-system agent prototype remains a preprint [45].	Scenario generation, simulation support, code generation, and natural-language interfaces; peer-reviewed evidence is available [46], while automated dispatch modeling is also reported in a preprint [47].	Application selection must reflect latency, safety criticality, uncertainty, and human oversight requirements.
Main strength and limitation	Strength: predictable and interpretable. Limitation: limited adaptability.	Strength: adaptive distributed optimization. Limitation: training instability and weak guarantees.	Strength: flexible reasoning and workflow orchestration. Limitation: limited grid-specific validation, hallucination, latency, tool-risk, and governance concerns.	Strength: knowledge accessibility and human-machine communication. Limitation: generated outputs may be inaccurate or physically infeasible.	No paradigm independently satisfies all grid requirements; verified hybrid deployment is the most defensible approach.

5. Key Application Domains

5.1. Voltage and Frequency Regulation

Voltage and frequency regulation are among the most operationally critical applications of multi-agent intelligence because renewable-rich grids contain geographically dispersed inverters, reduced synchronous inertia, variable generation, and rapidly changing power flows. MAS architectures can assign local controllers to distributed generators, storage systems, flexible loads, and microgrids so that active and reactive power support is coordinated using selected state information. Peer-reviewed studies have formulated secondary voltage regulation as a cooperative MARL problem [48], developed multi-agent deep meta-reinforcement learning for load-frequency control across operating conditions [43], and separated voltage restoration from current sharing through a two-stage MARL architecture for DC microgrids [42]. These studies demonstrate the value of decentralized inference, but practical deployment requires explicit stability margins, bounded actions, communication-delay tolerance, and protection-compatible response times. Evaluation should therefore include voltage deviation, frequency nadir, rate of change of frequency, settling time, overshoot, current-sharing error, control effort, constraint violations, and resilience under topology or communication changes.

5.2. Distributed Energy Management, Demand Response, and Active Grid Peripheries

Distributed energy management requires coordinated scheduling of renewable generators, storage devices, flexible loads, microgrids, aggregators, and market transactions across multiple timescales. Hierarchical trust-region MARL has been used for collaborative multi-energy microgrid scheduling [44], while federated deep reinforcement learning has enabled coordinated policies across differently sized microgrids without centralizing complete local operating data [49]. On the demand side, load-aggregator agents and consensus-based communication have coordinated time-of-use and direct-load-control resources

under uncertain device availability [50]. These applications increasingly extend to the active grid periphery, where prosumers, rooftop photovoltaics, batteries, electric vehicles, smart buildings, and low-voltage network devices create bidirectional power flows. Event-driven edge agents can reduce unnecessary communication while maintaining control performance [51], and recent systematic evidence identifies hybrid consensus, distributed optimization, MARL, and edge intelligence as important techniques for peripheral control [52]. Performance should be assessed using operating cost, renewable utilization, peak reduction, prosumer welfare, carbon emissions, energy not served, communication burden, computation time, privacy preservation, and physical-constraint compliance.

5.3. Electric Mobility, Fault Management, and Self-Healing Operation

Electric vehicles and self-healing systems illustrate how MAS can convert distributed assets from passive grid disturbances into coordinated flexibility and resilience resources. Grid-aware MARL has been used to accept, delay, or redirect charging demand while considering local network capacity [53]; multi-agent scheduling has combined traffic forecasting, dynamic pricing, and renewable self-consumption for highway charging stations [54]; and heterogeneous graph MARL has coordinated routing and charging across coupled transportation-power networks [55]. Other work models electric vehicles, charging stations, renewable resources, storage, and the grid as interacting agents for coordinated charging and bidirectional exchange [56]. Following network failures, centralized training with decentralized execution and invalid-action masking has produced constraint-compliant switching sequences for distributed load restoration [57], while MARL with dynamic network reconfiguration has restored prioritized loads under changing topology [58]. A recent review synthesizes multi-agent fault detection, localization, isolation, reconfiguration, restoration, and post-event voltage control [59]. Evaluation should therefore include charging cost, station utilization, waiting time, departure state of charge, battery degradation, restored-load percentage, restoration time, switching actions, unserved energy, and fault-classification accuracy.

Table 2. Multi-agent System Techniques, Contributions, Evaluation Metrics, and Challenges by Application Domain

Application domain	Principal MAS techniques	Main contribution to grid operation	Core evaluation metrics	Critical challenges	Representative evidence
Voltage control	Cooperative actor-critic MARL; attention-based communication; consensus control; distributed secondary control; safety-filtered policies	Coordinates inverter reactive power, storage, tap changers, and distributed generators to restore bus voltages while limiting local control conflicts	Maximum and mean voltage deviation; voltage-violation rate; settling time; overshoot; reactive-power use; network losses	Guaranteeing closed-loop stability; measurement noise; communication delay; controller interaction; transfer across feeder topologies	[42, 48]
Frequency control	Multi-agent deep meta-RL; distributed policy-gradient control; adaptive droop agents; cooperative load-frequency control	Allocates active-power and reserve responses among areas, storage systems, controllable loads, and renewable generators under low-inertia	Frequency nadir; rate of change of frequency; steady-state error; settling time; tie-line deviation; regulation cost	Fast response requirements; rare-event training; uncertain inertia; cyberattacks; policy instability under unseen disturbances	[43]

		conditions			
Distributed energy management and scheduling	Hierarchical MARL; federated RL; distributed optimization; actor–critic scheduling; cooperative game agents	Coordinates multi-energy resources, storage, interconnected microgrids, energy trading, and day-ahead or real-time dispatch	Operating cost; renewable curtailment; storage degradation; carbon emissions; import dependency; convergence time; energy balance	Large joint action spaces; competing owners; non-stationarity; forecast errors; data heterogeneity; physical constraint enforcement	[44, 49]
Demand response and load management	Aggregator agents; consensus algorithms; cooperative MARL; price-responsive agents; direct-load-control agents	Aggregates flexible appliances, buildings, industrial loads, and prosumers to reduce peaks, follow dispatch signals, and provide ancillary services	Peak reduction; peak-to-average ratio; energy cost; rebound peak; comfort deviation; participation rate; response accuracy	Consumer privacy; uncertain availability; fairness; preference modelling; rebound effects; incentive compatibility	[50]
Electric-vehicle coordination and V2G	Grid-aware MADDPG; graph MARL; multi-station transformer MARL; routing–charging agents; bidirectional V2G agents	Coordinates charging location, time, rate, redirection, and discharge to reduce congestion, absorb renewables, and provide grid flexibility	Charging cost; station utilization; waiting time; transformer loading; voltage violations; departure state of charge; battery aging	Mobility uncertainty; driver acceptance; battery degradation; heterogeneous chargers; cybersecurity; transport–grid coupling	[53, 54, 55, 56]
Fault detection and self-healing	Protection agents; MARL switching agents; invalid-action masking; distributed diagnosis; dynamic reconfiguration; restoration negotiation	Detects, isolates, and localizes faults, forms networked microgrids, restores priority loads, and coordinates available distributed resources	Fault-detection accuracy; isolation time; restored-load percentage; restoration duration; unserved energy; switching count; violations	Incomplete fault data; miscoordination with protection; combinatorial switching spaces; damaged communications; cyber deception	[57, 58, 59]
Active grid periphery management	Event-driven edge agents; consensus MAS; distributed MPC; ADMM; local markets; blockchain-enabled agents; hybrid MARL	Manages low-voltage networks, prosumers, rooftop solar, batteries, smart buildings, and localized flexibility near the consumer boundary	Response latency; communication volume; hosting capacity; local voltage quality; curtailment; transaction efficiency; edge-computing load	Intermittent connectivity; limited edge resources; asymmetric information; interoperability; device diversity; cyber-physical uncertainty	[51, 52]

Table 2 shows that application suitability is determined primarily by operational timescale, physical risk, and validation requirements. Voltage and frequency services require fast inference, stability assurance, and formal safety constraints; scheduling and demand response permit slower economic coordination; electric-vehicle applications require joint modeling of mobility, network capacity, and battery health; and self-healing requires constraint-compliant decisions under damaged or incomplete information. A layered architecture combining deterministic local control, distributed MAS coordination, adaptive MARL, and operator-supervised restoration is therefore more credible than unrestricted end-to-end autonomy.

6. Critical Challenges and Research Gaps

6.1. Scalable Coordination under Heterogeneous Agent and Communication Conditions

Scalability is not simply a matter of increasing agent count; it also involves heterogeneous observations, action spaces, objectives, devices, and response times. As distributed energy resources proliferate, centralized training faces expanding joint state-action spaces, non-stationarity, and costly retraining. A peer-reviewed benchmark study highlights the need for evaluation across grid sizes and topologies [60], while topology generalization remains an unresolved issue in a graph-reinforcement-learning survey that is still a preprint [61]. Strong scalability has been reported for residential microgrid coordination [62], but such results remain sensitive to reward structure, information assumptions, and training distribution. Communication is another deployment constraint because many MARL models assume reliable and timely message exchange; a recent preprint specifically examines robustness to delay, packet loss, perturbation, and bandwidth limits [63]. Future work should prioritize sparse and event-triggered communication, asynchronous learning, graph-based neighborhood aggregation, agent clustering, and topology-invariant policies tested under changing populations, failures, and communication quality.

6.2. Cybersecure and Interoperable Multi-agent Ecosystems

Autonomous grid agents enlarge the cyber-physical attack surface because adversaries may manipulate measurements, rewards, model updates, messages, identities, prompts, memory, or executable tools. A recent smart-grid cybersecurity survey available as a preprint catalogs false-data injection, denial-of-service, malware, and compromised-device threats [64]. Peer-reviewed work has investigated adversarial training and federated heterogeneous learning for attack-resilient scheduling [65], while 6G-enabled grid architectures introduce additional edge, slicing, virtualization, and AI-specific vulnerabilities [66]. Security is also weakened by fragmented standards and incompatible information models. Semantic interoperability using large language models can translate among standards such as IEEE SEP 2.1, IEC 61968, and OpenADR, but generated translations require deterministic validation [67]. Research should therefore integrate authenticated agent identities, least-privilege permissions, secure aggregation, attack-aware control, protocol certification, standardized ontologies, and interoperable test profiles across vendors and grid layers.

6.3. Explainable Autonomy and Operator Trust

Operator trust remains a barrier because technically accurate policies may still be rejected when their reasoning, uncertainty, and operational consequences are unclear. Interpretable fuzzy systems can preserve predictive performance while exposing demand relationships [68], and explainability has been integrated with resilient, battery-aware microgrid control [69]. A U.S. national-laboratory technical report argues that generative AI for grid operations should support transparent human-AI collaboration, model validation, and certification [70], while a peer-reviewed systematic review positions human-in-the-loop governance as an important condition for safe energy-sector autonomy [71]. Future research should develop multi-agent explanations that reveal local and collective contributions, feasible alternatives, uncertainty, violated constraints, and consequences. Trust should be evaluated through operator studies, decision time, calibration,

override quality, workload, and post-incident auditability rather than accuracy alone.

6.4. Evidence Maturity and Preprint Sensitivity

The preprint-exclusion sensitivity analysis did not materially change the principal conclusions for traditional MAS and MARL. Peer-reviewed journal and conference evidence independently supports voltage and frequency control, distributed scheduling, demand response, electric-vehicle coordination, and service restoration [42, 43, 44, 48, 49, 50, 53, 54, 57, 58]. Accordingly, the central conclusion that layered MAS/MARL architectures can improve distributed coordination in simulation, benchmark, and selected implementation settings does not depend on preprint evidence. The same analysis also leaves intact the conclusions concerning the need for physical constraints, deterministic safety layers, cybersecurity controls, interoperability, and human oversight. Delay-robust Volt/Var control has also been explored in a preprint [72], but that source is not required for these core conclusions.

The effect is different for Agentic AI. Peer-reviewed grid-specific evidence now exists for customized large-model agents and review-level frameworks [22, 26, 27], but several influential sources used to characterize general agent planning, multi-agent collaboration, power-system agent prototypes, and automated LLM-based dispatch remain preprints [24, 25, 45, 47, 73–75]. When these preprints are removed, claims about broad autonomous task decomposition, tool orchestration, and cross-functional grid workflows become substantially less certain. The evidence is therefore not mature enough to justify treating Agentic AI as a deployment-ready control paradigm on the same footing as established MAS or MARL. Its current role is better framed as an emerging supervisory, analytical, and human-interface layer that requires rigorous benchmarking, permission boundaries, engineering-solver validation, and field evaluation.

6.5. Regulatory, Market, Economic, Workforce, and Data-Governance Barriers

Regulatory acceptance is likely to be a gating condition for autonomous grid agents. In the European Union, the AI Act classifies AI used as a safety component in the management or operation of critical infrastructure, including electricity supply, as high risk, bringing requirements for risk management, data governance, technical documentation, logging, human oversight, accuracy, robustness, and cybersecurity into scope [76]. In North America, AI-enabled grid functions must also fit within existing reliability and cybersecurity obligations. The NERC Critical Infrastructure Protection standards address security management, personnel and training, electronic security perimeters, system security, incident response, configuration and vulnerability management, information protection, and supply-chain risk management [77]. These frameworks favor bounded autonomy, auditable interfaces, deterministic validation, controlled change management, and clear allocation of responsibility rather than unrestricted self-directed control.

Market design determines whether technically capable multi-agent coordination can create realizable value. FERC Order No. 2222 illustrates the importance of aggregation rules, minimum participation thresholds, locational constraints, metering and telemetry, bidding parameters, settlement, and coordination among regional operators, distribution utilities, aggregators, and retail regulators [78]. Comparable issues arise in other jurisdictions. A multi-agent controller can optimize flexibility in simulation yet remain economically ineffective if market products do not reward the service, baseline and settlement rules are misaligned, or distribution-level constraints prevent delivery. Economic evaluation should therefore extend beyond operating-cost reductions to lifecycle total cost of ownership, including communications, edge/cloud computing, data engineering, cybersecurity, software assurance, retraining, model monitoring, licensing, compliance, integration with legacy EMS/SCADA systems, operator training, insurance, and contingency support. Business cases should report net present value, payback period, avoided curtailment or outage cost, sensitivity to market prices and participation rates, and the cost of maintaining safe fallback operation.

Workforce implications are similarly socio-technical rather than purely algorithmic. Human-in-the-loop

evidence in the energy sector indicates that operators will increasingly supervise, validate, and override AI-assisted decisions rather than simply execute manually derived schedules [71]. An agency report on AI for energy also emphasizes both near-term grid opportunities and the need for safe, secure, and responsible deployment [79]. Utilities therefore need training in model limitations, uncertainty interpretation, cyber incident response, data stewardship, validation of agent-generated actions, and escalation procedures. Automation should be designed to avoid deskilling and automation bias: operators must retain the competence and authority to recognize unsafe recommendations, revert to verified control modes, and manage degraded communications or model failure.

Data governance connects regulatory compliance, cybersecurity, privacy, and economic viability. Multi-agent platforms require explicit rules for data ownership, provenance, quality, purpose limitation, retention, access control, sharing, and deletion; they also need model and prompt versioning, action logs, and traceable links between source data and operational recommendations. Household demand, electric-vehicle mobility, and prosumer data may contain personal or commercially sensitive information, while operational network data can be security sensitive. Federated learning reduces raw-data centralization but does not by itself guarantee privacy, security, or lawful secondary use. Deployment architectures should therefore separate control-critical and personal data domains, apply least-privilege access, record lineage and consent or contractual authority where applicable, validate synthetic-data provenance, and define which data may be transmitted to external generative models. These governance controls should be specified before model training and procurement rather than added after technical integration.

7. Future Research Directions

7.1. Converged Digital Twin, Edge Intelligence, and Blockchain Infrastructure

Future MAS platforms should integrate digital twins, edge intelligence, and distributed trust mechanisms selectively rather than treating them as universally beneficial add-ons. Digital twins can provide synchronized grid representations for state estimation, contingency rehearsal, policy testing, and predictive maintenance [80, 81], while prosumer-level studies illustrate their use in energy optimization [82]. Edge nodes can execute latency-sensitive inference near substations, inverters, buildings, and charging assets, whereas blockchain is better suited to identity, provenance, and settlement than to fast inner-loop control [83]. Research should prioritize semantic interoperability, real-time twin calibration, edge-cloud workload partitioning, privacy-preserving data exchange, failure-aware synchronization, and lifecycle maintenance. Evaluation should include latency, fidelity, energy overhead, cybersecurity, scalability, recoverability, and total cost of ownership.

7.2. Generative Scenario Intelligence and Physics-Constrained Synthetic Data

Generative AI can expand scenario intelligence by producing renewable trajectories, demand profiles, contingencies, cyber events, and rare operating states that are poorly represented in historical datasets. Peer-reviewed work demonstrates statistically validated synthetic injections for transmission-grid learning [84] and synthetic samples for power-system fault classification [85]. Physics-informed diffusion models have also been proposed for feasible and diverse power-flow-state generation, although the cited study remains a preprint [86]. Digital-twin and generative-AI integration is emerging as a broader pathway for forecasting, warning, and dispatch [87]. Future studies should move beyond visual or statistical realism toward physics consistency, temporal coherence, topology awareness, privacy protection, and coverage of cascading or adversarial events. Synthetic datasets should include provenance records, quality scores, uncertainty bounds, and train-on-synthetic/test-on-real validation, and generated scenarios should be screened by power-flow, stability, and protection models before operational use.

7.3. Trustworthy Agentic AI, Benchmarking, and Certification Readiness

Future Agentic AI frameworks should be designed for assurance rather than autonomy alone. Peer-reviewed power-system work identifies tool misuse, cascading misinformation, model ambiguity, and weak human oversight as important barriers [27], while a recent trustworthiness framework used for additional conceptual coverage remains a preprint [73]. Lifecycle risk management and evaluation principles are also available in the NIST Generative AI Profile [88]. For grid use, these principles should be translated into explicit agent contracts, permission boundaries, safety shields, uncertainty reports, action logs, verified tool interfaces, and operator escalation for high-consequence decisions. Certification readiness requires benchmarks spanning multiple networks, renewable conditions, communication failures, cyberattacks, and distribution shifts. ISO/IEC 42001:2023 provides an AI-management-system framework [89], ISO/IEC 25059:2023 defines quality characteristics for AI systems [90], ISO/IEC 24029-2:2023 addresses robustness assessment using formal methods [91], and ISO/IEC 42006:2025 specifies requirements for bodies auditing and certifying AI management systems [92]. Future work should align grid benchmarks with these governance structures and develop graded assurance levels based on operational scope, autonomy, explainability, resilience, cybersecurity, and worst-case safety.

8. Conclusion and Limitations

This review shows that multi-agent artificial intelligence provides a useful architectural foundation for distributed smart-grid coordination, but the maturity of the evidence varies substantially by paradigm. Traditional MAS offers interpretable rule- and protocol-based coordination, while MARL provides adaptive distributed optimization under uncertainty. Peer-reviewed evidence supports applications in voltage and frequency regulation, distributed scheduling, demand response, electric-vehicle coordination, self-healing, and active grid management, although most studies remain simulation based and difficult to compare because test systems, reward functions, communication assumptions, and metrics vary. Generative AI can support scenario generation, simulation, documentation, and operator interaction. Agentic AI is less mature: peer-reviewed grid-specific work is emerging, but several influential architecture and workflow claims still rely on preprints or early prototypes. The preprint-exclusion sensitivity analysis therefore leaves the MAS/MARL conclusions largely unchanged while narrowing claims about deployment-ready Agentic AI. Across paradigms, the most defensible pathway is a layered system combining deterministic local control, distributed intelligence, validated power-system solvers, cybersecurity and data-governance controls, market-compatible incentives, and human oversight. Future research should prioritize reproducible benchmarks, hardware-in-the-loop and field evaluation, lifecycle economic analysis, workforce readiness, interoperable standards, auditable governance, and staged certification before unrestricted autonomous deployment in critical power systems.

Appendix A. Database-Specific Search Strategies

The search concepts were translated to the syntax supported by each source. Straight quotation marks were used for phrase searching, Boolean operators were capitalized, and database-native date filters were applied where available. The eligible publication window was 1 January 2023 through 31 August 2026; English-language eligibility was applied through database filters where supported and otherwise during screening. Records were exported with available bibliographic metadata and DOI information for de-duplication. Because ScienceDirect Advanced Search permits only a limited number of Boolean operators in a single query, that source was searched in three broader, syntactically valid batches and the operational-function criterion was applied during screening. This platform-specific adaptation avoids presenting a non-executable Boolean string as though it were reproducible.

Scopus. TITLE-ABS-KEY((“multi-agent system” OR “multi-agent systems” OR “multiagent system” OR “multiagent systems” OR “multi-agent reinforcement learning” OR MARL OR “Agentic AI” OR “LLM agent” OR “LLM agents” OR “large language model agent” OR “large language model agents” OR (“generative AI” AND agent)) AND (“smart grid” OR “smart grids” OR microgrid OR microgrids OR “power system” OR “power systems” OR “distribution network” OR “distribution networks” OR “distributed energy resource” OR “distributed energy resources” OR DER OR “renewable energy”) AND (control OR coordination OR dispatch OR scheduling OR “energy management” OR “demand response” OR restoration OR “voltage control” OR “frequency control” OR trading)) AND PUBYEAR > 2022 AND PUBYEAR < 2027 AND (LIMIT-TO(LANGUAGE, “English”)).

Web of Science Core Collection. TS=((“multi-agent system” OR “multi-agent systems” OR “multiagent system” OR “multiagent systems” OR “multi-agent reinforcement learning” OR MARL OR “Agentic AI” OR “LLM agent” OR “LLM agents” OR “large language model agent” OR “large language model agents” OR (“generative AI” AND agent)) AND (“smart grid” OR “smart grids” OR microgrid OR microgrids OR “power system” OR “power systems” OR “distribution network” OR “distribution networks” OR “distributed energy resource” OR “distributed energy resources” OR DER OR “renewable energy”) AND (control OR coordination OR dispatch OR scheduling OR “energy management” OR “demand response” OR restoration OR “voltage control” OR “frequency control” OR trading)) AND PY=(2023-2026) AND LA=(English).

IEEE Xplore. The full cross-platform expression was entered in Command Search (not in the structured Advanced Search fields), because Command Search supports nested Boolean expressions. Operators were entered in uppercase. Results were restricted to publication years 2023-2026; English-language eligibility was verified during screening. Query: (“multi-agent system” OR “multi-agent systems” OR “multiagent system” OR “multiagent systems” OR “multi-agent reinforcement learning” OR MARL OR “Agentic AI” OR “LLM agent” OR “LLM agents” OR “large language model agent” OR “large language model agents” OR (“generative AI” AND agent)) AND (“smart grid” OR “smart grids” OR microgrid OR microgrids OR “power system” OR “power systems” OR “distribution network” OR “distribution networks” OR “distributed energy resource” OR “distributed energy resources” OR DER OR “renewable energy”) AND (control OR coordination OR dispatch OR scheduling OR “energy management” OR “demand response” OR restoration OR “voltage control” OR “frequency control” OR trading).

ScienceDirect, batch 1 (Title, abstract or keywords). (“multi-agent system” OR “multi-agent reinforcement learning” OR MARL) AND (“smart grid” OR microgrid OR “power system” OR “distribution network” OR “distributed energy resource” OR “renewable energy”).

ScienceDirect, batch 2 (Title, abstract or keywords). (“Agentic AI” OR “LLM agent”) AND (“smart grid” OR microgrid OR “power system” OR “distribution network” OR “distributed energy resource” OR “renewable energy”).

ScienceDirect, batch 3 (Title, abstract or keywords). (“large language model agent” OR (“generative AI” AND agent)) AND (“smart grid” OR microgrid OR “power system” OR “distribution network” OR “distributed energy resource” OR “renewable energy”).

SpringerLink supplementary coverage check. Keyword searches used (“multi-agent” OR MARL OR “Agentic AI”) AND (“smart grid” OR microgrid OR “power system”), with publication years restricted to 2023-2026. Records already captured in the four core databases were not counted as independent PRISMA identifications.

ACM Digital Library supplementary coverage check. Advanced Search used Boolean operators in uppercase: (“multi-agent” OR MARL OR “Agentic AI”) AND (“smart grid” OR microgrid OR “power system”), with a custom publication-date range of 2023-2026. Overlapping records were used for coverage verification or full-text retrieval rather than as a separate identification stream.

Google Scholar supplementary search. (“multi-agent reinforcement learning” OR “Agentic AI” OR “multi-

agent system") ("smart grid" OR microgrid OR "power system") (control OR coordination OR "energy management" OR restoration). Google Scholar was used only for supplementary discovery and citation chasing; records entering screening were de-duplicated against the core database set.

Conflict of Interest

The author declares no conflict of interest.

Author Contributions

Larbi Kwadwo conceptualized the review, developed the search strategy, conducted the literature search and study selection, extracted and synthesized the evidence, designed the comparative taxonomy and analytical tables, interpreted the findings, prepared the original manuscript, and revised the paper critically for important intellectual content. The author read and approved the final version of the manuscript.

Acknowledgment

The author acknowledges the contributions of researchers whose scholarly works provided the evidence synthesized in this review.

References

[1] International Energy Agency. (2024). *Renewables 2024: Analysis and Forecasts to 2030*. IEA.

[2] International Energy Agency. (2025). *Renewables 2025*. IEA.

[3] International Energy Agency. (2023). *Electricity grids and secure energy transitions*. IEA.

[4] Jendoubi, I., & Bouffard, F. (2023). Multi-agent hierarchical reinforcement learning for energy management. *Applied Energy, 332*, 120500. doi: 10.1016/j.apenergy.2022.120500

[5] Monfaredi, F., Shayeghi, H., & Siano, P. (2023). Multi-agent deep reinforcement learning-based optimal energy management for grid-connected multiple energy carrier microgrids. International Journal of *Electrical Power and Energy Systems, 153*, 109292. doi: 10.1016/j.ijepes.2023.109292

[6] Joshi, A., Capezza, S., Alhaji, A., & Chow, M. Y. (2023). Survey on AI and machine learning techniques for microgrid energy management systems. *IEEE/CAA Journal of Automatica Sinica, 10(7)*, 1513–1529. doi: 10.1109/JAS.2023.123657

[7] Xie, J., Ajagekar, A., & You, F. (2023). Multi-agent attention-based deep reinforcement learning for demand response in grid-responsive buildings. *Applied Energy, 342*, 121162. doi: 10.1016/j.apenergy.2023.121162

[8] Nuvvula, R. S. S., Kumar, P. P., Ahammed, S. R., Satyanarayana, V., Babu, B. H., Reddy, R. S. S., & Ali, A. (2024). Distributed multi-agent reinforcement learning for autonomous management of renewable energy microgrids. *Proceedings of the 2024 12th International Conference on Smart Grid* (pp. 298–304). IEEE. doi: 10.1109/icSmartGrid61824.2024.10578150

[9] Wang, Z., Xiao, F., Ran, Y., Li, Y., & Xu, Y. (2024). Scalable energy management approach of residential hybrid energy system using multi-agent deep reinforcement learning. *Applied Energy, 367*, 123414. doi: 10.1016/j.apenergy.2024.123414

[10] Wang, X., Liu, S., Xu, Q., & Shao, X. (2025). Distributed multi-agent reinforcement learning for multi-objective optimal dispatch of microgrids. ISA *Transactions, 158*, 130–140. doi: 10.1016/j.isatra.2025.01.009

[11] Hribar, J., Mohorčič, M., & Čampa, A. (2025). Improving energy autonomy of positive energy districts using multi-agent deep reinforcement learning. *Scientific Reports, 15*, 27798. doi: 10.1038/s41598-025-12554-x

[12] Bashyal, A., Boroukhian, T., Veerachanchai, P., Naransukh, M., & Wicaksono, H. (2025). Multi-agent deep reinforcement learning based demand response and energy management for heavy industries with discrete manufacturing systems. *Applied Energy, 392*, 125990. doi: 10.1016/j.apenergy.2025.125990

[13] Khan, M. M., Ahmad, S., Naeem, M., Ali, M. U., & Kim, S. H. (2026). A unified multi-agent optimization framework for intelligent PV-integrated smart energy systems. *Scientific Reports, 16*, 20039. doi: 10.1038/s41598-026-50395-4

[14] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., *et al.* (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372(71).* doi: 10.1136/bmj.n71

[15] Rethlefsen, M. L., Kirtley, S., Waffenschmidt, S., Ayala, A. P., Moher, D., Page, M. J., et al. (2021). PRISMA-S: An extension to the PRISMA statement for reporting literature searches in systematic reviews. *Systematic Reviews, 10*, 39. doi: 10.1186/s13643-020-01542-z

[16] Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large language model-based multi-agents: A survey of progress and challenges. *Proceedings of International Joint Conference on Artificial Intelligence* (pp. 8048–8057). doi: 10.24963/ijcai.2024/890

[17] Chen, S., Liu, Y., Han, W., Zhang, W., & Liu, T. (2024). A survey on LLM-based multi-agent systems: Recent advances and new frontiers in application. arXiv preprint, arXiv.2412.17481.

[18] Maldonado, D., Cruz, E., Abad Torres, J., Cruz, P. J., & Gamboa Benitez, S. del P. (2024). Multi-agent systems: A survey about their components, framework, and workflow. *IEEE Access, 12*, 80950–80975. doi: 10.1109/ACCESS.2024.3409051

[19] Dragomir, O. E., & Dragomir, F. (2025). A decentralized hierarchical multi-agent framework for smart grid sustainable energy management. *Sustainability, 17(12)*, 5423. doi: 10.3390/su17125423

[20] Piccialli, F., Chiaro, D., Sarwar, S., Cerciello, D., Qi, P., & Mele, V. (2025). AgentAI: A comprehensive survey on autonomous agents in distributed artificial intelligence for Industry 4.0. *Expert Systems with Applications, 291*, 128404. doi: 10.1016/j.eswa.2025.128404

[21] Arévalo, P., Ochoa-Correa, D., Villa-Ávila, E., Iñiguez-Morán, V., & Astudillo-Salinas, P. (2025). Systematic review of hierarchical and multi-agent optimization strategies for peer-to-peer energy management and electric machines in microgrids. *Applied Sciences, 15(9)*, 4817. doi: 10.3390/app15094817

[22] Kiasari, M., & Aly, H. H. (2026). Agentic artificial intelligence for smart grids: A comprehensive review of autonomous, safe, and explainable control frameworks. *Energies, 19(3)*, 617. doi: 10.3390/en19030617

[23] Liu, D., Ren, F., Yan, J., Su, G., Gu, W., & Kato, S. (2024). Scaling up multi-agent reinforcement learning: An extensive survey on scalability issues. *IEEE Access, 12*, 94610–94631. doi: 10.1109/ACCESS.2024.3410318

[24] Huang, X., Liu, W., Chen, X., Wang, X., Wang, H., Lian, D., Wang, Y., Tang, R., & Chen, E. (2024). Understanding the planning of large language model agents: A survey. arXiv preprint, arXiv.2402.02716.

[25] Masterman, T., Besen, S., Sawtell, M., & Chao, A. (2024). The landscape of emerging AI agent architectures for reasoning, planning, and tool calling: A survey. arXiv preprint, arXiv.2404.11584.

[26] Dong, H., Yang, H., Miao, Y., Zhu, J., Cheung, K., Hua, H., Sun, C., Li, S., Wang, Z., *et al.* (2026). Operating smart grids by customizing large model agents. *Communications Engineering, 5*, 114. doi: 10.1038/s44172-026-00709-1

[27] Ghosh, S., & Mittal, G. (2026). Agentic AI systems in electrical power systems engineering: Current state-of-the-art and challenges. *Frontiers in Artificial Intelligence, 9*, 1814651. doi: 10.3389/frai.2026.1814651

[28] Billah, M., Zeb, K., Uddin, W., Imran, M. A., Alatawi, K. S., Almasoudi, F. M., & Khalid, M. (2025). Decentralized multi-agent control for optimal energy management of neighborhood-based hybrid microgrids in real-time networking. *Results in Engineering, 27*, 106337. doi:

10.1016/j.rineng.2025.106337

[29] Aeggegn, D. B., Nyakoe, G. N., & Wekesa, C. (2024). A state-of-the-art review on energy management techniques and optimal sizing of distributed energy resources in grid-connected multi-microgrids. *Cogent Engineering, 11(1)*, 2340306. doi: 10.1080/23311916.2024.2340306

[30] Brahmia, I., & Amrani, A. (2025). Hierarchical smart energy management framework based on multi-agent reinforcement learning and distributed model predictive control for peer-to-peer energy-sharing communities. *Proceedings of the 8th International Conference on Electrical Engineering and Green Energy* (pp. 372–377). IEEE. doi: 10.1109/CEEGE66246.2025.11182327

[31] Das, O., Zafar, M. H., Sanfilippo, F., Rudra, S., & Kolhe, M. L. (2024). Advancements in digital twin technology and machine learning for energy systems: A comprehensive review of applications in smart grids, renewable energy, and electric vehicle optimisation. *Energy Conversion and Management: X, 24*, 100715. doi: 10.1016/j.ecmx.2024.100715

[32] Djebali, S., Guérard, G., & Taleb, I. (2024). Survey and insights on digital twins design and smart grid applications. *Future Generation Computer Systems, 153*, 234–248. doi: 10.1016/j.future.2023.11.033

[33] Yıldırım, F., Yalman, Y., Bayındır, K. Ç., & Terciyanlı, E. (2025). Comprehensive review of edge computing for power systems: State of the art, architecture, and applications. *Applied Sciences, 15(8)*, 4592. doi: 10.3390/app15084592

[34] Syamala, M., Gowri, U., Babu, D. V., Nisha, S. A., Ahmed, M. A., & Muniyandy, E. (2024). Transactive energy management system for smart grids using multi-agent modelling and blockchain. *Sustainable Computing: Informatics and Systems, 43*, 101001. doi: 10.1016/j.suscom.2024.101001

[35] Ma, H., Zhang, H., Tian, D., Yue, D., & Hancke, G. P. (2024). Optimal demand response-based dynamic pricing strategy via multi-agent federated twin-delayed deep deterministic policy gradient algorithm. *Engineering Applications of Artificial Intelligence, 133*, 108012. doi: 10.1016/j.engappai.2024.108012

[36] Abdellatif, A. A., Shaban, K., & Massoud, A. M. (2025). Blockchain-enabled distributed learning for enhanced smart grid security and efficiency. *Computers & Electrical Engineering, 123*, 110012. doi: 10.1016/j.compeleceng.2024.110012

[37] Izmirlioglu, Y., Pham, L., Son, T. C., & Pontelli, E. (2024). A survey of multi-agent systems for smartgrids. *Energies, 17(15)*, 3620. doi: 10.3390/en17153620

[38] Qu, Y., Ma, J., & Wu, F. (2024). Safety constrained multi-agent reinforcement learning for active voltage control. *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence* (pp. 184–192). doi: 10.24963/ijcai.2024/21

[39] Shi, H., Fang, L., Chen, X., Gu, C., Ma, K., Zhang, X., Zhang, Z., Gu, J., & Lim, E. G. (2024). Review of the opportunities and challenges to accelerate mass-scale application of smart grids with large-language models. *IET Smart Grid, 7(6)*, 737–759. doi: 10.1049/stg2.12191

[40] Cheng, Y., Zhao, H., Zhou, X., Zhao, J., Cao, Y., Yang, C., & Cai, X. (2025). A large language model for advanced power dispatch. *Scientific Reports, 15*, 8925. doi: 10.1038/s41598-025-91940-x

[41] Henao, F., Edgell, R., Sharma, A., & Olney, J. (2025). AI in power systems: A systematic review of key matters of concern. *Energy Informatics, 8*, 76. doi: 10.1186/s42162-025-00529-1

[42] Li, F., Tu, W., Zhou, Y., Li, H., Zhou, F., Liu, W., & Hu, C. (2025). Distributed secondary control for DC microgrids using two-stage multi-agent reinforcement learning. *International Journal of Electrical Power and Energy Systems, 164*, 110335. doi: 10.1016/j.ijepes.2024.110335

[43] Li, J., Dai, J., & Cui, H. (2025). Bionic cooperative load frequency control in interconnected grids: A multi-agent deep meta-reinforcement learning approach. *Applied Energy, 379*, 124906. doi: 10.1016/j.apenergy.2024.124906

[44] Xu, X., Xu, K., Zeng, Z., Tang, J., He, Y., Shi, G., & Zhang, T. (2024). Collaborative optimization of multi-energy

multi-microgrid system: A hierarchical trust-region multi-agent reinforcement learning approach. *Applied Energy, 375*, 123923. doi: 10.1016/j.apenergy.2024.123923

[45] Jin, H., Kim, K., & Kwon, J. (2025). GridMind: LLMs-powered agents for power system analysis and operations. arXiv preprint, arXiv.2509.02494.

[46] Jia, M., Cui, Z., & Hug, G. (2025). Enhancing LLMs for power system simulations: A feedback-driven multi-agent framework. *IEEE Transactions on Smart Grid, 16(6)*, 5556–5572. doi: 10.1109/TSG.2025.3589114

[47] Yang, X., Lin, C., Yang, Y., Wang, Q., Liu, H., Hua, H., & Wu, W. (2025). Large language model-powered automated modelling and optimization of active distribution network dispatch problems. arXiv preprint, arXiv.2507.21162.

[48] Wang, T., Ma, S., Tang, Z., Xiang, T., Mu, C., & Jin, Y. (2023). A multi-agent reinforcement learning method for cooperative secondary voltage control of microgrids. *Energies, 16(15)*, 5653. doi: 10.3390/en16155653

[49] Zhang, Y., Ren, Y., Liu, Z., Li, H., Jiang, H., Xue, Y., Ou, J., Hu, R., Zhang, J., & Gao, D. W. (2025). Federated deep reinforcement learning for varying-scale multi-energy microgrids energy management considering comprehensive security. *Applied Energy, 380*, 125072. doi: 10.1016/j.apenergy.2024.125072

[50] Zeng, P., Xu, J., & Zhu, M. (2024). Demand response strategy based on the multi-agent system and multiple-load participation. *Sustainability, 16(2)*, 902. doi: 10.3390/su16020902

[51] Zhang, X., Liu, Y., Gu, S., Tian, Y., & Gao, Y. (2025). Event-driven edge agent framework for distributed control in distribution networks. *Energies, 18(11)*, 2734. doi: 10.3390/en18112734

[52] Mamun, S., Ioannou, S., Christofides, N. G., & Darwish, M. (2026). Multi-agent systems for decentralized control and management of active power grid peripheries: A systematic review. *Applied Sciences, 16(14)*, 6863. doi: 10.3390/app16146863

[53] Yang, Y., & Tang, R. (2026). Cooperative multi-agent reinforcement learning for grid-aware EV charging management with cross-site redirection. *Energy and AI, 23*, 100664. doi: 10.1016/j.egyai.2025.100664

[54] Zhou, J., Xiang, Y., Zhang, X., Sun, Z., Liu, X., & Liu, J. (2025). Optimal self-consumption scheduling of highway electric vehicle charging station based on multi-agent deep reinforcement learning. *Renewable Energy, 238*, 121982. doi: 10.1016/j.renene.2024.121982

[55] Mei, Z., Zhang, Y., Jiang, H., Xue, Y., Zhang, J., & Gao, D. W. (2026). Multi-agent heterogeneous graph reinforcement learning for electric vehicle routing and charging scheduling in coupled power-transportation networks. *Applied Energy, 403*, 126958. doi: 10.1016/j.apenergy.2025.126958

[56] Li, Y., Zhao, H., Song, Z., Chong, X., Chen, R., Li, L., & Wang, X. (2026). Coordinated interaction of electric vehicles, charging stations, and the power grid using multi-agent reinforcement learning. *Advances in Applied Energy, 22*, 100274. doi: 10.1016/j.adapen.2026.100274

[57] Vu, L., Vu, T., Vu, T. L., & Srivastava, A. (2024). Multi-agent deep reinforcement learning for distributed load restoration. *IEEE Transactions on Smart Grid, 15(2)*, 1749–1760. doi: 10.1109/TSG.2023.3310893

[58] Si, R., Chen, S., Zhang, J., Xu, J., & Zhang, L. (2024). A multi-agent reinforcement learning method for distribution system restoration considering dynamic network reconfiguration. *Applied Energy, 372,* 123625. doi: 10.1016/j.apenergy.2024.123625

[59] Hamdeen, I., Badran, E. A., Kotb, M. F., & Elgamal, M. (2025). Self-healing multi-agent techniques in electric power distribution systems: A review. *Renewable and Sustainable Energy Reviews, 224*, 116132. doi: 10.1016/j.rser.2025.116132

[60] Bogucki, D. J., Lepak, Ł., Parashar, S., Błachowski, B., & Wawrzyński, P. (2025). EnEnv 1.0: Energy grid environment for multi-agent reinforcement learning benchmarking. *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems* (pp. 361–370).

[61] Hassouna, M., Holzhüter, C., Lytaev, P., Thomas, J., Sick, B., & Scholz, C. (2024). Graph reinforcement

learning for power grids: A comprehensive survey. arXiv preprint, arXiv.2407.04522.

[62] Wang, C., Wang, M., Wang, A., Zhang, X., Zhang, J., Ma, H., Yang, N., Zhao, Z., Lai, C. S., & Lai, L. L. (2025). Multiagent deep reinforcement learning-based cooperative optimal operation with strong scalability for residential microgrid clusters. *Energy, 314,* 134165. doi: 10.1016/j.energy.2024.134165

[63] Liu, Z., Li, Y., Wang, J., Tu, J., Hong, Y., Li, F., Liu, Y., Sugawara, T., & Tang, Y. (2025). Robust and efficient communication in multi-agent reinforcement learning. arXiv preprint, arXiv.2511.11393.

[64] Mashima, D., Chen, Y., Roomi, M. M., Lakshminarayana, S., & Chen, D. (2024). Cybersecurity for modern smart grids against emerging threats. arXiv preprint, arXiv.2404.04466.

[65] Qi, H., Ma, W., & Zhao, B. (2025). Federated heterogeneous multi-agent deep reinforcement learning-based attack resilience scheduling for heterogeneous multi-integrated energy systems. *Energy Reports, 13*. doi: 10.1016/j.egyr.2025.05.085

[66] Sanjalawe, Y. K., Al-E'mari, S., Fraihat, S., Makhadmeh, N., & Alzubi, E. (2025). AI-powered smart grids in the 6G era: A comprehensive survey on security and intelligent energy systems. *IEEE Open Journal of the Communications Society, 6*, 7677–7719. doi: 10.1109/OJCOMS.2025.3609144

[67] Fatemi, N., & Fattahi, J. (2025). Adaptable semantic interoperability in heterogeneous smart grids using large language models. *Energy Reports, 14*, 5774–5789. doi: 10.1016/j.egyr.2025.12.011

[68] Alfayan, M., & Hagras, H. (2025). Towards an explainable artificial intelligence approach for smart grid systems. *Discover Artificial Intelligence*. doi: 10.1007/s44163-025-00261-5

[69] Nejati Amiri, M. H., & Guéniat, F. (2026). Explainable deep reinforcement learning for resilient and battery-aware microgrid control. *Energy Conversion and Management, 353*, 121215. doi: 10.1016/j.enconman.2026.121215

[70] Choi, S. L., Jain, R., Feng, C., Emami, P., Zhang, H., Hong, J., Kim, T., Park, S., Ding, F., Baggu, M., & Kroposki, B. (2024). Generative AI for power grid operations (NREL/TP-5D00-91176). *National Renewable Energy Laboratory*. doi: 10.2172/2477920

[71] Nguyen, A. T., & Ahn, Y. (2026). Human-in-the-loop artificial intelligence in the energy sector: A systematic review of paradigm shifts toward Industry 5.0. *Advances in Applied Energy, 22*, 100267. doi: 10.1016/j.adapen.2026.100267

[72] Cheng, H., Luo, H., Liu, Z., Sun, W., Li, W., & Li, Q. (2024). Reinforcement learning based robust Volt/Var control in active distribution networks with imprecisely known delay. arXiv preprint, arXiv.2402.17268.

[73] Al-Refai, O., Shahbaz, I., Husseinat, A. A., Mandulak, M., Kim, J., & Hammad, E. (2026). Engineering trustworthy Agentic AI for critical systems. arXiv preprint, arXiv.2607.18548.

[74] Ren, J., Gao, H., Zhao, L., Kang, Q., Xiao, G., & Sun, Y. (2026). Multi-agent reinforcement learning for low-carbon peer-to-peer energy trading among self-interested microgrids. arXiv preprint, arXiv.2604.08973.

[75] Tran, K. T., Dao, D., Nguyen, M. D., Pham, Q. V., O'Sullivan, B., & Nguyen, H. D. (2025). Multi-agent collaboration mechanisms: A survey of large language models. arXiv preprint, arXiv.2501.06322.

[76] European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689. Retrieved from https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[77] North American Electric Reliability Corporation. (2026). Critical Infrastructure Protection (CIP) reliability standards. Retrieved from https://www.nerc.com/standards/reliability-standards/cip

[78] Federal Energy Regulatory Commission. (2020). Participation of distributed energy resource aggregations in markets operated by regional transmission organizations and independent system operators, Order No. 2222, 172 FERC ¶ 61,247.

[79] U.S. Department of Energy. (2024). *AI for energy: Opportunities for a modern grid and clean energy economy. Office of Critical and Emerging Technologies*. Retrieved from

https://www.energy.gov/cet/articles/ai-energy

[80] Al-Shetwi, A. Q., Atawi, I. E., El-Hameed, M. A., & Abuelrub, A. (2025). Digital twin technology for renewable energy, smart grids, energy storage and vehicle-to-grid integration: Advancements, applications, key players, challenges and future perspectives in modernising sustainable grids. *IET Smart Grid, 8(1)*, e70026. doi: 10.1049/stg2.70026

[81] Alharbey, R., Shafiq, A., Daud, A., Dawood, H., Bukhari, A., & Alshemaimri, B. (2024). Digital twin technology for enhanced smart grid performance: Integrating sustainability, security, and efficiency. *Frontiers in Energy Research, 12*, 1397748. doi: 10.3389/fenrg.2024.1397748

[82] Ali, S. M., Khan, B., & Ullah, Z. (2025). Deep learning-enabled digital twins for prosumers: A holistic energy management framework for smart grids using deep reinforcement learning and big data analytics. *International Journal of Energy Research, 2025*, 6618907. doi: 10.1155/er/6618907

[83] Zahid, H., Zulfiqar, A., Adnan, M., Iqbal, M. S., Shah, A., Abbasi, U., & Mohamed, S. E. G. (2025). Transforming nano grids to smart grid 3.0: AI, digital twins, blockchain, and the metaverse revolutionizing the energy ecosystem. *Results in Engineering, 27*, 105850. doi: 10.1016/j.rineng.2025.105850

[84] Gillioz, M., Dubuis, G., & Jacquod, P. (2025). A large synthetic dataset for machine learning applications in power transmission grids. *Scientific Data, 12*, 168. doi: 10.1038/s41597-025-04479-x

[85] Shukla, P. K., & Deepa, K. (2025). AI-based synthetic data generation techniques for improved fault classification in power systems. *Ain Shams Engineering Journal, 16(8)*, 103485. doi: 10.1016/j.asej.2025.103485

[86] Wang, J., Upadhyay, D., Zaman, M., & Srikantha, P. (2025). Synthetic power-flow data generation using physics-informed denoising diffusion probabilistic models. arXiv preprint, arXiv.2504.17210.

[87] Wu, Q., & Zhao, W. (2026). Review of digital twin and generative artificial intelligence for operation and control of new-type power systems. *Cyber-Physical Energy Systems, 2(3)*, 22–36. doi: 10.1016/j.cpes.2026.06.003

[88] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). *National Institute of Standards and Technology*. doi: 10.6028/NIST.AI.600-1

[89] International Organization for Standardization and International Electrotechnical Commission. (2023). ISO/IEC 42001:2023 information technology: Artificial intelligence management system.

[90] International Organization for Standardization and International Electrotechnical Commission. (2023). ISO/IEC 25059:2023 software engineering: Systems and software quality requirements and evaluation, quality model for AI systems.

[91] International Organization for Standardization and International Electrotechnical Commission. (2023a). ISO/IEC 24029-2:2023 artificial intelligence: Assessment of the robustness of neural networks, Part 2: Methodology for the use of formal methods.

[92] International Organization for Standardization and International Electrotechnical Commission. (2025). ISO/IEC 42006:2025 information technology: Requirements for bodies providing audit and certification of artificial intelligence management systems.

Copyright © 2026 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (CC BY 4.0).