*energies*

MDPI

*Review*

# The Cognitive Power Mini-Grid with Distributed AI, Semantic Control and Agentic Autonomy: Concepts, Applications, Challenges and Future Directions

**Iacovos Ioannou [1,2,3,*] and Saher Javaid [4,5]**

[1] Department of Computer Science, Neapolis University Paphos, 8042 Paphos, Cyprus
[2] Department of Computer Science, University of Cyprus, 1678 Nicosia, Cyprus
[3] CYENS Centre of Excellence, 1016 Nicosia, Cyprus
[4] Department of Information Engineering, Kanazawa Gakuin University, Kanazawa 920-1392, Japan; saher@kanazawa-gu.ac.jp
[5] Graduate School of Advanced Science and Technology, Japan Advanced Institute of Science and Technology, Nomi 923-1292, Japan
* Correspondence: iakovos.ioannou@nup.ac.cy

## Abstract

Power mini-grids are being transformed from locally automated electrical systems into cyber-physical ecosystems in which heterogeneous distributed energy resources, storage, converters, flexible demand and uncertain external conditions must be coordinated. Existing surveys commonly treat microgrid control, artificial intelligence (AI), multi-agent systems, communications, digital twins and cybersecurity as separate research streams. In this survey, the cognitive power mini-grid is introduced as a unifying paradigm in which physical and social contexts are perceived by distributed agents, task-relevant semantic information is exchanged, auditable coordination readiness is checked and proposed actions are subject to independent safety checks. A corpus of 247 scholarly publications, standards and technical sources is synthesized across microgrid engineering, distributed AI, semantic communication, language models, federated learning, neuro-symbolic AI, digital twins, causal reasoning, runtime assurance, cybersecurity and community energy markets. A six-layer architecture is proposed in which fast deterministic control is separated from semantic coordination, distributed learning, agentic deliberation and assurance. Representative studies are compared by problem, method, control horizon, information assumptions, validation environment, hardware-in-the-loop (HIL) status, reported outcome, limitations and architectural relevance. Applications are organized into balancing, resilience, demand-response, maintenance and inter-mini-grid markets. Cross-layer challenges are identified in stability under asynchronous interaction, semantic interoperability, hallucination, common-knowledge failure, edge resources, privacy and cyber-physical security constraints. A research agenda is developed around proof-carrying actions, causal digital twins, edge small language models, federated multimodal foundation models, neuromorphic semantic control and human-agent governance. Cognition is therefore positioned as a safety-gated coordination capability above verified physical control loops rather than as a replacement for established control.

**Keywords:** power mini-grid; microgrid; distributed artificial intelligence; multi-agent systems; semantic communication; semantic control; common knowledge; agentic AI; large language models; digital twins; trustworthy AI; federated learning



Academic Editors: Cong Zhang, Qin Wang, Zipeng Liang and Xin Yin

Received: 5 September 2026
Revised: 11 September 2026
Accepted: 16 September 2026
Published: 19 September 2026

## 1. Introduction

Power mini-grids combine distributed energy resources (DERs), energy storage, controllable loads and a defined electrical boundary that permits grid-connected or islanded operation. Their practical importance is increasing because renewable generation can be integrated, resilience can be improved, remote electrification can be supported and local energy resources can be managed by communities. Nevertheless, high shares of converter-interfaced generation, uncertain demand, mobile storage, climate-driven disturbances and market participation make their operation more complex than the centralized dispatch assumptions embedded in traditional energy management systems (EMSs) presuppose. The mature microgrid literature has consequently developed a layered body of control methods. The microgrid was first defined as a cluster of loads and microsources that appears to the wider system as a single controllable entity and can island and reconnect autonomously [1]. Early overviews consolidated the operating concept [2] and the management functions required for grid-connected and islanded operation [3]. Hierarchical primary, secondary and tertiary control was then formalized for droop-controlled converters [4] and the responsibilities of centralized and distributed controllers at each level were systematized [5]. A subsequent task-force review traced the trend toward advanced energy management, model predictive control (MPC) and communication-supported distributed coordination [6]. The control strategies, stabilization techniques and power architectures of direct-current (DC) microgrids were reviewed in a two-part survey [7,8], which complements the alternating-current (AC) systems addressed by the earlier works.

Artificial intelligence (AI) has added forecasting, anomaly detection, adaptive scheduling and model-free decision-making to this foundation. Multi-agent systems (MASs) distribute decisions across generators, storage devices, loads and supervisory entities. The concepts, engineering practice and standards of power-system MASs were consolidated in a two-part survey [9,10]. An early multi-agent microgrid controller demonstrated market-based local dispatch [11], agent-based DER management with demand response was developed for intelligent microgrids [12] and the resulting body of multi-agent microgrid control was surveyed by Kantamneni et al. [13]. Consensus theory supplies the mechanism through which such agents agree on shared quantities over a communication graph [14]. Reinforcement learning (RL) and multi-agent RL (MARL) learn policies for economic operation under uncertainty and multi-agent common-knowledge RL (MACKRL) has shown, outside the power domain, that conditioning group policies on what agents provably share improves coordination [15]. Federated learning (FL) supports cross-site learning without centralizing raw measurements [16], although its open problems in privacy, robustness and heterogeneity remain substantial [17]. Graph learning exploits electrical topology [18]. These developments are significant, but in the assessment of the present survey, the field remains fragmented. Most studies optimize a single task, assume a fixed communication model, treat data as semantically uniform and place limited emphasis on what each agent knows about the knowledge of other agents.

Semantic and task-oriented communication offers a different abstraction. Instead of reconstruction fidelity being maximized for every transmitted measurement, the information useful for a control or inference task is optimized. In a mini-grid, the required message may be "reserve unavailable for the next peak" rather than a high-rate stream of battery states, provided that the receiver can interpret the message, assess its validity horizon and map it to a safe action. The theoretical foundation is supplied by several lines of work. A theory of semantic communication that measures information by meaning rather than by symbols was outlined by Bao et al. [19], semantic and goal-oriented communication was positioned as a defining feature of sixth-generation (6G) networks [20] and the term itself was clarified for machine-intelligence settings [21]. Context, semantics and task-oriented

transmission were unified in a tutorial treatment [22], an end-to-end learned semantic transmission system was demonstrated for text [23] and task-oriented communication was formulated as a design principle for 6G [24]. Age of information provides the timeliness measure with which the freshness of a received state is quantified [25] and control under communication constraints establishes the minimum information rate at which a plant can be stabilized [26]. In the microgrid domain, neuromorphic event-driven semantic communication has demonstrated that sparse, state-relevant exchanges can be embedded in converter-dominated systems [27].

Large language models (LLMs), small language models (SLMs) and tool-using agents create a further transition from prediction to deliberation. The Transformer architecture on which these models are built [28] and the foundation-model paradigm of pretraining once and adapting to many tasks [29] supply the technical basis. An agent can interpret operator goals, retrieve grid rules, call power-flow or optimization tools, construct a multi-step plan and explain the proposed action. Retrieval-augmented generation (RAG) grounds generated text in retrieved documents [30], ReAct interleaves reasoning steps with tool actions [31] and Toolformer trains a model to decide when to call external tools [32]. Emerging grid–agent systems already combine semantic reasoning with numerical solvers and sandboxed validation. Grid-Agent pairs a planning agent with a validation agent that executes proposals in a sandbox with rollback [33] and X-GridAgent organizes planning, coordination and action layers around grid-analysis tools [34]. A recent review of agentic AI for smart grids catalogs such systems together with their safety and explainability requirements [35]. However, these systems are early prototypes and should not be interpreted as evidence that unrestricted LLM control is suitable for safety-critical infrastructure. The central engineering question is therefore not whether an LLM can issue a command but how deliberative agents can be separated from millisecond control, grounded in verified state, constrained by physical and organizational rules and prevented from acting when evidence is stale or incomplete.

That question is addressed in this survey through the concept of the cognitive power mini-grid. Eight coordinated capabilities are included in the definition of cognition. The first four are (i) perception of multimodal state; (ii) abstraction of observations into task-relevant semantics; (iii) maintenance of local, shared and common knowledge; and (iv) reasoning over goals and constraints. The remaining capabilities are (v) use of domain tools and learned policies, (vi) prediction of consequences in a digital twin, (vii) certification of actions before execution and (viii) learning from outcomes under explicit governance. Cognition is thereby distinguished from generic automation and from a single AI algorithm. The four conditions under which the composition of these capabilities is cognitive rather than merely integrated are stated in Section 3.1.

The novelty of the survey is fourfold. First, distributed AI, semantic control, common-knowledge reasoning and agentic autonomy are unified in one mini-grid architecture. Second, communication semantics and epistemic state are treated as control variables rather than as background infrastructure. Third, proof-carrying agent actions are introduced as a concrete assurance pattern in which each high-level action is accompanied by assumptions, constraints, predicted consequences, validity horizon and a machine-checkable certificate. Fourth, a benchmark and deployment roadmap is proposed in which physical, communication, intelligence and governance dimensions are evaluated together.

The evolution of the proposed paradigm is summarized in Figure 1. The remainder of the paper is organized as follows. The review scope and analytical framework are defined in Section 2. The cognitive mini-grid concepts and unified architecture are developed in Section 3. The enabling research streams are examined in depth and compared across studies in Section 4. Operational applications and open challenges are reviewed in

Sections 5 and 6. Future directions, evaluation requirements and the deployment roadmap are presented in Section 7. The survey is concluded in Section 8.

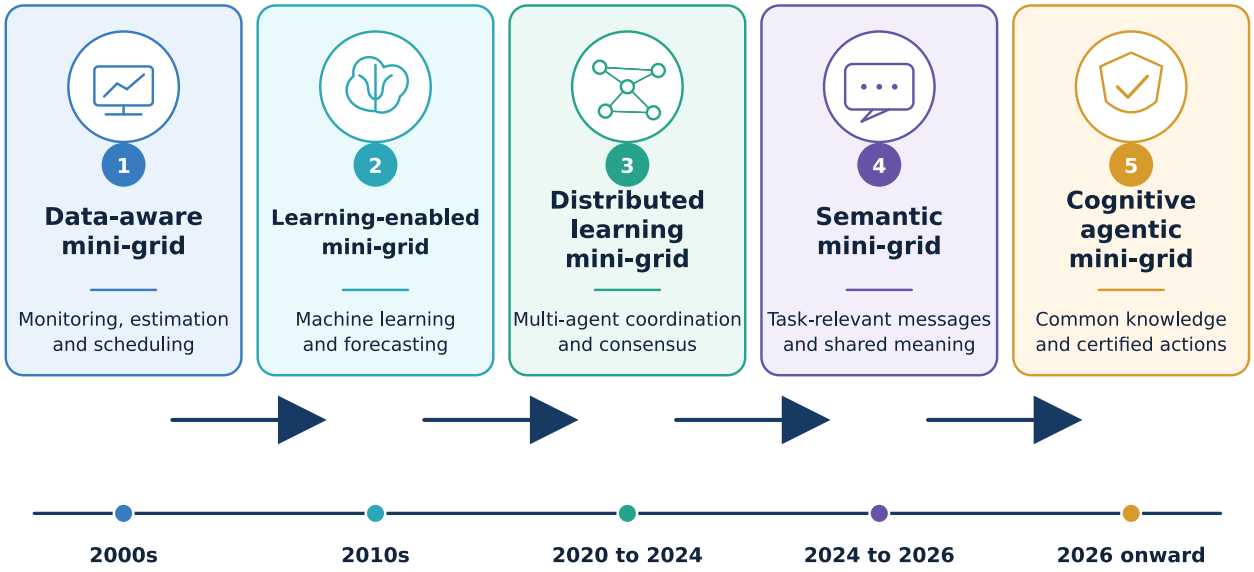


**Figure 1.** Conceptual progression from data-aware automation to cognitive agentic mini-grid autonomy, with overlapping illustrative periods and cognition augmenting established control.

The material is arranged so that the core argument can be separated from the reference material that supports it. The core argument is carried by four elements: the six-layer architecture and its interface contracts, the unified cross-family comparison, the proof-carrying action pattern and the evaluation requirements that follow from them. These elements are developed in Sections 3–7 and are summarized, with pointers to the evidence that supports each of them, in the table of principal findings stated in Section 2.4. The remaining material is supporting reference content of three kinds: the per-cluster examination tables, which record what each examined study did and the limits of its evidence; the reporting checklists, which state what a future study must disclose; and the standards and roadmap material, which places the architecture in the existing normative landscape. A reader interested only in the principal findings may therefore read Sections 3 and 4 together with the principal findings stated in Section 2.4 and consult the checklists when a specific study or evaluation is designed. The survey is deliberately comprehensive in the reference material because no single existing review spans the seven research streams that the architecture composes and a shorter treatment would reintroduce the fragmentation that the survey is intended to remove.

## 2. Review Scope and Analytical Framework

In this section, the methodological basis of the survey is established. The evidence base, source categories, temporal coverage and terminology are defined; the search strings, the databases, the screening stages, the inclusion and exclusion criteria, the composition of the retained corpus and the treatment of preprints are stated; the screening and classification procedure is given as an algorithm; the four research questions and the physical, information, intelligence and governance assessment dimensions are formulated; and the survey is positioned relative to the existing review traditions. A common evidence synthesis and comparison protocol is then specified; a unified ordinal scale for horizon, authority, evidence, system scale and communication assumptions is defined, together with the rules under which studies may be compared at all; the principal findings of the survey are stated, along with the evidence on which each of them rests; and a compact analytical

model connecting physical state, semantic encoding, agent beliefs, high-level actions and safety-governed execution is introduced. Because heterogeneous systems, datasets, control horizons and outcome measures are covered, a structured narrative synthesis is adopted instead of a statistical meta-analysis.

### *2.1. Scope, Source Types and Temporal Coverage*

The evidence base covers archival journal and conference publications, books, standards, technical reports and selected preprints available through August 2026. Peer-reviewed work is prioritized for established claims. Preprints are used only to identify emerging directions such as grid foundation models, LLM-based power-system agents, federated foundation models and neuro-symbolic grid control; their status is explicitly distinguished from deployment evidence. The bibliography contains 247 entries and spans seven clusters: microgrid foundations, distributed intelligence, AI and learning, semantic communication, agentic AI, trustworthy autonomy and standards/deployment.

A mini-grid is treated as a bounded distribution-level energy system that may be grid-connected, islanded or remote. This usage follows the microgrid architectures and control functions described in the reference textbook of Hatziargyriou [36] and the reviews that summarize the state of microgrid research [37]; its technologies and key drivers [38]; and its architecture, policy and future trends [39]. The community-scale and remote-electrification meaning of mini-grid follows the agency reports on renewable mini-grids [40] and on mini-grid markets for remote electrification [41]. In this survey, microgrid is used when terminology in the cited literature is discussed and mini-grid is used when community-scale or remote deployment is emphasized. AC and DC microgrids with DERs have been compared [42] and AC, DC and hybrid microgrid technologies have been reviewed [43]. The cognitive architecture proposed here is intended to apply to all three configurations, but it does not replace converter control, protection, state estimation or grid codes.

Source discovery, screening and inclusion followed the protocol stated below. Because the protocol was not prospectively registered and no quantitative study-selection flow was maintained, the work is classified as a structured narrative survey rather than a Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) systematic review.

The search strategy is stated explicitly so that the coverage of the survey can be judged and repeated. Nine query groups were executed, each combining a domain term with a technology term. The domain terms were microgrid, mini-grid, nanogrid, distributed energy resource and distribution feeder. The technology terms were, per group, (i) hierarchical control, droop control and secondary restoration; (ii) distributed optimization, consensus, alternating direction method of multipliers and distributed model predictive control; (iii) multi-agent system, agent-based control and belief–desire–intention agent; (iv) reinforcement learning, deep reinforcement learning, multi-agent reinforcement learning and energy management; (v) semantic communication, task-oriented communication, goal-oriented communication and age of information; (vi) large language model, language agent, tool use, retrieval-augmented generation and agentic artificial intelligence; (vii) digital twin, co-simulation and hardware-in-the-loop; (viii) federated learning, foundation model, graph neural network and continual learning; and (ix) runtime assurance, neuro-symbolic reasoning, explainable artificial intelligence, cybersecurity, false-data injection and certification. Each group was executed against IEEE Xplore, Scopus, Web of Science and Google Scholar and, additionally, against arXiv for the sixth to ninth groups. Standards and normative documents were retrieved directly from the IEEE, IEC, NIST and ISO repositories rather than through the bibliographic databases because their identifiers rather than their titles are the reliable search key. Backward and forward citation chasing was applied to every

source that was retained after full-text assessment, which is the mechanism through which the foundational literature of the field was reached.

Four screening stages were applied. At identification, records returned by the query groups were merged and duplicates were removed. At title and abstract screening, records without a traceable connection to a bounded distribution-level energy system or to one of the four research questions were removed. At full-text assessment, each remaining record was read against the inclusion and exclusion criteria of Table 1 and the evidence template stated in Section 2.4; a record was retained when it satisfied IC2 and IC6 together with at least one of IC1, IC4 or IC5, satisfied IC3 whenever it reported an evaluation and met none of EC1 to EC6; IC1, IC4 and IC5 describe alternative kinds of contribution and therefore cannot be required jointly. At inclusion, the record was assigned to one of the seven evidence clusters and to the layer of the architecture to which its contribution applies. Screening and full-text assessment were carried out by the authors. Because no prospective register of the counts excluded at each stage was maintained, the composition of the retained corpus is reported in full in Table 2 so that its balance can be inspected directly.

**Table 1.** Inclusion and exclusion criteria applied at full-text assessment.

| Code | Criterion |
|---|---|
| *Inclusion* | |
| IC1 | A method, architecture, protocol or analytical result is contributed that applies to a bounded distribution-level energy system or is a domain-independent result on which such a method is built. |
| IC2 | The contribution addresses at least one of the four research questions and can be assigned to at least one layer of the architecture. |
| IC3 | For a record that reports an evaluation, the validation environment, the exchanged information and the produced action are reported in sufficient detail for the evidence template to be completed; items that are not reported are recorded as unavailable rather than inferred. |
| IC4 | A foundational definition, a normative requirement or a standard is provided, on which later work in the corpus depends. |
| IC5 | A benchmark, dataset, simulator or test feeder is provided that is relevant to the evaluation of such methods. |
| IC6 | The record is written in English and its bibliographic metadata are verifiable at the level of the publisher, the standards body or the preprint server. |
| *Exclusion* | |
| EC1 | The record is a duplicate, an extended abstract, a poster or a presentation without an accompanying full text. |
| EC2 | The artificial-intelligence contribution has no traceable relevance to power or energy systems and is not a foundational result used by a retained source. |
| EC3 | The electrical contribution concerns transmission-level or generation-fleet problems only and does not transfer to a bounded distribution-level system. |
| EC4 | An empirical performance claim cannot be attributed to a stated method, plant model or validation environment; a conceptual or normative source is not excluded by this criterion because it reports no experiment. |
| EC5 | The bibliographic metadata could not be verified, or the venue could not be identified. |
| EC6 | The record repeats, without addition, a contribution that is already represented in the corpus by an archival version of the same work. |

Rows in italics are group headings.

The search, screening and classification procedure is stated as Algorithm 1 so that it can be repeated and its recorded counts are stated in the paragraph that follows the algorithm and in Table 2. Lines 1 to 6 of the algorithm are the identification stage: every combination of a domain term and a technology term of the nine groups is executed against the four bibliographic databases, with arXiv added for the sixth to ninth groups; the standards are retrieved from their repositories by identifier; and the union is deduplicated. Lines 7 to 10 are the title and abstract screening, in which a record is kept only when it concerns

a bounded distribution-level energy system or supplies a foundational result, a standard or an evaluation resource on which such a system depends, and addresses at least one of the four research questions. Lines 11 to 14 are the full-text assessment, in which a record is retained only when it satisfies IC2 and IC6, together with at least one of IC1, IC4 or IC5; satisfies IC3 whenever it reports an evaluation; and meets none of EC1 to EC6, with the evidence template of Section 2.4 completed for every retained record. Lines 15 to 19 are the citation-chasing loop, which submits the backward and forward citations of every retained record to the same screening and assessment and repeats until no new record is retained; this loop is the mechanism through which the foundational literature of the field was reached. Lines 20 to 23 are the classification, in which every retained record is assigned to the evidence cluster of its primary contribution, to the layers of the architecture to which it contributes and to the research questions it addresses, with every examined study placed on the unified scale defined in Section 2.4. Line 24 is the metadata verification of 31 August 2026, at which point the preprints that had reached archival publication were replaced by their archival versions, with line 25 returning the retained corpus.

**Table 2.** Composition of the retained evidence base by source type and period of publication.

| Source Type | Up to 2010 | 2011 to 2015 | 2016 to 2020 | 2021 to 2026 | Total |
|---|---|---|---|---|---|
| Journal article | 24 | 38 | 51 | 43 | 156 |
| Conference paper | 6 | 3 | 23 | 7 | 39 |
| Book or book chapter | 6 | 4 | 3 | 1 | 14 |
| Standard, technical report or dataset | 1 | 2 | 11 | 11 | 25 |
| Preprint | 0 | 0 | 2 | 11 | 13 |
| **Total** | **37** | **47** | **90** | **73** | **247** |

The bold values in the table represent the column totals.

The counts of $R_0$ after line 6, of $R_1$ after line 10 and of the records excluded at line 13 were not recorded because the protocol was not prospectively registered, so what was recorded is reported instead: the criteria applied at each stage, which are those of Table 1 and of lines 9 and 13 of the algorithm; the retained corpus by source type and by period, which is given in Table 2; and the assignment of the retained entries to the seven evidence clusters, namely 41 entries on microgrid foundations, 41 on distributed intelligence, 42 on AI and learning, 25 on semantic communication, 21 on agentic AI, 39 on trustworthy autonomy and 38 on standards and deployment. Each entry is counted once, in the cluster of its primary contribution, whereas its assignment to the layers of the architecture may be multiple. The procedure and these counts are a record of how the retained corpus was formed rather than a study-selection flow in the sense of a systematic review and they are provided so that the search can be repeated and the balance of the evidence base can be inspected.

The temporal distribution of the retained sources is reported in Table 2. Sources published from 2011 onward form the core of the corpus because hierarchical microgrid control was consolidated in that period and every technology stream examined here developed after it. Earlier sources were retained only when they supplied a definition, a theorem or a standard on which the later corpus depends, which is why the corpus reaches back to the mathematical theory of communication and its semantic reading. The metadata cutoff was 31 August 2026. Of the 247 retained entries, 73 were published from 2021 onward and 90 between 2016 and 2020 so that 163 entries, that is, two-thirds of the corpus, post-date 2015, while the remaining third secures the definitions, the theorems and the standards that the recent work assumes.

**Algorithm 1** Search, screening and classification procedure of the survey.

**Require:** domain terms $D$; technology-term groups $G_1, \ldots, G_9$; bibliographic databases $B$; standards repositories $S$; inclusion criteria IC1 to IC6 and exclusion criteria EC1 to EC6 (Table 1); research questions (RQs) RQ1 to RQ4; evidence clusters $K_1, \ldots, K_7$; layers $L_1, \ldots, L_6$
**Ensure:** retained corpus $R$ with cluster, layer, research-question and class assignments
1: $R_0 \leftarrow \emptyset$ ▷ identification
2: **for** $i = 1, \ldots, 9$ **do**
3: $B_i \leftarrow B \cup \{\text{arXiv}\}$ **if** $i \geq 6$, **otherwise** $B_i \leftarrow B$
4: $R_0 \leftarrow R_0 \cup \bigcup_{d \in D,\, t \in G_i} \textsc{Query}(d \wedge t,\, B_i)$
5: **end for**
6: $R_0 \leftarrow \textsc{Deduplicate}\big(R_0 \cup \textsc{RetrieveByIdentifier}(S)\big)$
7: $R_1 \leftarrow \emptyset$ ▷ title and abstract screening
8: **for** each $r \in R_0$ **do**
9: **if** $r$ concerns a bounded distribution-level energy system, or supplies a foundational result, a standard or an evaluation resource on which such a system depends, **and** addresses at least one $\text{RQ}_q$ **then** $R_1 \leftarrow R_1 \cup \{r\}$
10: **end for**
11: $R_2 \leftarrow \emptyset$ ▷ full-text assessment
12: **for** each $r \in R_1$ **do**
13: **if** $r$ satisfies IC2 and IC6 and at least one of IC1, IC4 or IC5, satisfies IC3 whenever it reports an evaluation **and** meets none of EC1 to EC6 **then** complete the evidence template of $r$ and set $R_2 \leftarrow R_2 \cup \{r\}$
14: **end for**
15: **repeat** ▷ citation chasing
16: $N \leftarrow \big(\textsc{Backward}(R_2) \cup \textsc{Forward}(R_2)\big) \setminus R_2$
17: $N \leftarrow$ the subset of $N$ retained by lines 8 to 14
18: $R_2 \leftarrow R_2 \cup N$
19: **until** $N = \emptyset$
20: **for** each $r \in R_2$ **do** ▷ classification
21: cluster$(r) \leftarrow$ the $K_k$ of the primary contribution of $r$; layers$(r) \leftarrow$ the subset of $\{L_1, \ldots, L_6\}$ to which $r$ contributes; questions$(r) \leftarrow$ the subset of $\{\text{RQ1}, \ldots, \text{RQ4}\}$ that $r$ addresses
22: **if** $r$ is an examined study **then** class$(r) \leftarrow (\text{T}, \text{A}, \text{E}, \text{S}, \text{C})$ on the unified scale of Section 2.4
23: **end for**
24: $\textsc{VerifyMetadata}(R_2)$; replace every preprint in $R_2$ that has reached archival publication by its archival version
25: **return** $R \leftarrow R_2$

The treatment of preprints is stated separately because the frontier streams of the survey are the streams in which archival publication lags furthest behind the work itself. Thirteen of the 247 retained entries, that is, 5.3 percent, are cited as preprints or author-posted manuscripts and they account for 51 of the 695 individual source citations outside the two listing tables, i.e., Table 3 and the table that records the class position of every examined study in Section 4, that is, 7.3 percent. Seven entries are represented by archival versions in the retained bibliography, namely, the multi-agent conversation framework [44], the meta-programming framework [45], the survey of language-model agents [46], the account of the third wave of neuro-symbolic artificial intelligence [47], the formulation of task-oriented communication for 6G [24], the survey of large language models [48] and the edge-LLM review published in the Association for Computing Machinery (ACM) journal, *ACM Computing Surveys*. Three rules govern the retained nonarchival versions. First, their publication status is distinguished from the validation environment: peer-reviewed background can support a general concept but cannot independently confirm a particular prototype's claims. The final column of Table 3 therefore gives related peer-reviewed background, not validation of the preprint. Second, outcomes attributed to a preprint are identified as source-reported findings and are not treated as independently confirmed results. Evidence classes record the environment actually reported, including E2 when closed-loop simulation is described, while publication status remains a separate qualification. Third, established definitions and normative requirements are grounded in

archival or official sources; preprints are used to describe emerging directions and their explicitly limited evaluations.

**Table 3.** Preprints and author-posted manuscripts (13 entries) retained after the metadata update, their roles and related peer-reviewed background.

| Preprint | Role in the Survey | Peer-Reviewed Anchor |
| --- | --- | --- |
| Proximal policy optimization [49] | Names one algorithm family in the enumeration of policy-gradient methods | Surveys of reinforcement learning for microgrid control and management [50,51] |
| Grid foundation models [52] | Establishes that foundation-model pretraining is being attempted for grid operation | Survey of artificial-intelligence techniques across smart-grid applications [53] |
| Trustworthiness layer for power-system foundation models [54] | Establishes that a trust layer is being proposed for such models | Artificial-intelligence risk-management framework and risk-management guidance [55,56] |
| Semantic communication for 6G [57] | Reports the open questions of the semantic-communication agenda | Tutorial treatment of context, semantics and task-oriented transmission [22] |
| General-purpose model report [58] | Describes the capability level assumed by the agentic literature | Transformer architecture and the survey of language-model agents [28,46] |
| Open-weight model family [59] | Establishes that open-weight models of a deployable size exist | Transformer architecture [28] |
| Foundation-model paradigm [29] | Names the pretrain-and-adapt paradigm and its risks | Transformer architecture [28] |
| Grid-Agent [33] | Describes a planning and validation agent prototype for grid operation | Peer-reviewed review of agentic artificial intelligence for smart grids [35] |
| X-GridAgent [34] | Describes a layered agentic prototype for grid analysis | Peer-reviewed review of agentic artificial intelligence for smart grids [35] |
| Introduction to smart-grid digital twins [60] | Describes an emerging smart-grid specialization of established digital-twin concepts | Digital twin in industry and its state of the art [61] |
| Neuro-symbolic voltage control [62] | Describes a neuro-symbolic control prototype for a distribution testbed | Account of the third wave of neuro-symbolic artificial intelligence [47] |
| Problems in artificial-intelligence safety [63] | Supplies the vocabulary of specification and robustness failures | Comprehensive survey of safe reinforcement learning [64] |
| Unsolved problems in machine-learning safety [65] | Extends the same catalog to monitoring and alignment | Comprehensive survey of safe reinforcement learning [64] |

The edge-LLM review [66] is cited in its archival journal version and is excluded from the 13 preprints listed here. Related peer-reviewed background on wireless edge intelligence is provided by Park et al. [67].

### *2.2. Research Questions*

The synthesis is organized around four research questions. RQ1 concerns concepts: Which technical layers and knowledge representations are required for a mini-grid to be described as cognitive? RQ2 concerns applications: Which operational tasks can benefit from semantic and agentic coordination without violating real-time control boundaries? RQ3 concerns challenges: Which physical, informational, organizational and cyber risks prevent safe deployment? RQ4 concerns future directions: Which research programs can turn isolated AI demonstrations into interoperable, certifiable and resource-efficient mini-grid autonomy?

These questions lead to four assessment dimensions. The physical dimension concerns stability, power quality, protection, efficiency and resilience. The information dimension concerns timeliness, uncertainty, semantic sufficiency and shared knowledge. The intelli-

gence dimension concerns learning, planning, tool use and adaptation. The governance dimension concerns authority, accountability, cybersecurity, privacy and human oversight.

### *2.3. Distinction from Existing Surveys*

Existing surveys provide strong coverage of the individual research streams. Hierarchical control structures were classified by Bidram and Davoudi [5] and the trends toward advanced energy management and communication-supported coordination were reviewed by Olivares et al. [6]. RL solutions for microgrid control and management were surveyed by Barbalho et al. [50] and RL methods for microgrid energy management were reviewed with an outlook on future prospects by Zia et al. [51]. AI techniques across smart-grid applications were surveyed by Omitaomu and Niu [53] and agentic AI for smart grids was reviewed by Kiasari and Aly [35]. Their primary units of analysis are normally controllers, algorithms or application classes. The present survey instead asks how these components are composed into a cognitive stack and how semantic and epistemic assumptions propagate into physical decisions.

Intelligence is examined in the reviewed traditions, but their analyses stop at different system boundaries, as shown in Table 4. Microgrid-control reviews normally reach the electrical plant and provide the strongest treatment of stability, hierarchy and constraint satisfaction, but communication is frequently modeled as a reliable numerical channel. Multi-agent and distributed-control reviews move the decision boundary toward local assets and explain scalability, negotiation and consensus, yet the messages exchanged by the agents are usually assumed to have a fixed and mutually understood meaning. AI and RL reviews add adaptation and uncertainty handling, but safety is often represented by reward shaping or scenario constraints rather than an independently verifiable execution boundary. Semantic-communication surveys reverse the emphasis: they analyze meaning and task utility in depth, although the downstream electrical consequences of semantic loss are usually outside their scope. Agentic-AI surveys extend reasoning, tool use and natural-language interaction, whereas digital-twin and trustworthy-AI surveys emphasize validation and monitoring. None of these traditions alone closes the complete loop from physical evidence to semantic interpretation, distributed reasoning, action certification and bounded execution.

The comparison by tradition is completed by a comparison with the individual reviews that are closest to the present one, which is given in Table 5. Two of them require particular attention because their scope overlaps most with the present survey. The review of agentic artificial intelligence for smart grids by Kiasari and Aly [35] catalogs agentic, multi-agent, reinforcement-learning, digital-twin and physics-based control architectures and proposes an evaluation framework built on stability, safety, interpretability, certification readiness and interoperability; it is the reference catalog for Layer 5 of the present architecture and it is used as such in Section 4. It does not, however, state the condition under which an agent action becomes admissible on the plant, the contract that an agent message must satisfy before it may be acted upon, or the epistemic condition under which several agents may commit a shared resource and it does not relate the evidence available for a method to the authority that the method may be granted. The semantic and holonic architecture of Howell et al. [68] organizes highly distributed resources into nested holons that share semantic models and it is the closest antecedent of Layers 3 and 4; it does not specify the fields that make a semantic message actionable, the test by which shared knowledge is established before a coordinated action or the separation between a proposal and its execution and its authors themselves identify operational ontologies and safety interfaces as prerequisites that remain to be supplied.

**Table 4.** Positioning of the present survey relative to representative review traditions.

| Review Tradition | Primary Focus | Typical Strength | Gap Addressed Here |
|---|---|---|---|
| Microgrid control | Primary, secondary and tertiary control; droop, MPC and stability | Mature physical models and time-scale separation | Limited treatment of semantic knowledge, LLM agents and proof-carrying decisions |
| Multi-agent and distributed control | Consensus, distributed optimization and decentralized resource management | Scalability and removal of single-point control | Weak modeling of common knowledge, meaning alignment and agent authority |
| AI/RL for microgrids | Forecasting, scheduling, energy management and adaptive policies | Model-free operation under uncertainty | Safety, communication semantics and deployment governance often treated separately |
| Semantic/task-oriented communication | Meaning-preserving or goal-oriented transmission | Communication efficiency tied to task performance | Few power-specific ontologies, stability proofs or multi-timescale deployments |
| Agentic AI for smart grids | LLM reasoning, tool use, planning and explainability | Flexible interaction with complex workflows | Mini-grid constraints, edge operation, common knowledge and certification need deeper treatment |
| Digital-twin and trustworthy AI | Simulation, prediction, explainable artificial intelligence and runtime safety | Validation and operator support | Limited integration into a full semantic-agentic control lifecycle |
| **This survey** | Cognitive mini-grid composition | Unified architecture linking physical, semantic, learning and assurance layers | Joint concepts, applications, challenges, metrics and roadmap |

The bold values in the table represent the present survey.

**Table 5.** Comparison with the individual reviews and architectural proposals closest to the present survey. Layer numbers refer to the six-layer architecture of Section 3.

| Review or Proposal | Unit of Analysis | Layers Covered | Left Unspecified | Added Here |
|---|---|---|---|---|
| Bidram and Davoudi [5] | Controller functions of the hierarchy | 1, 2 | Communication meaning, agent authority, learning | Contracts I1 and I2; realization of the hierarchy in Layers 1 and 2 without change |
| Olivares et al. [6] | Control trends and energy management | 1, 2, 4 | Semantic and epistemic assumptions of coordination | Message contract; establishment test for common knowledge |
| Kantamneni et al. [13] | Multi-agent microgrid control | 2, 4 | Meaning of exchanged messages; admissibility of actions | Contracts I3 and I4; authority classes A1 to A5 |

**Table 5.** *Cont.*

| Review or Proposal | Unit of Analysis | Layers Covered | Left Unspecified | Added Here |
|---|---|---|---|---|
| Howell et al. [68] | Semantic and holonic organization | 3, 4 | Actionable message fields; establishment test; proposal-execution separation | Requirements one to three of the list that follows this table; Layers 5 and 6 |
| Omitaomu and Niu [53] | AI techniques by application | 4, 5 | Evidence needed before authority; interface to control | Coupling rule; evidence class per study (Section 4) |
| Gündüz et al. [22] | Semantic and task-oriented transmission | 3 | Electrical consequence of semantic loss; authority | Contract I3; safety preservation after compression as an evidence item |
| Barbalho et al. [50]; Zia et al. [51] | Reinforcement-learning methods for microgrids | 4 | Safe-set evidence, communication assumptions and authority | Classes C and E applied per study; A3 cap for simulation-only methods |
| Kiasari and Aly [35] | Agentic architectures and evaluation criteria | 4, 5 | Admissibility condition; message contract; epistemic precondition; evidence–authority relation | Requirements one to six of the list that follows this table; certificate acceptance conditions |
| Billanes et al. [69] | Resilience strategies and performance metrics | 1, 4 | Cognitive coordination; certification of restoration actions | Physical resources as the bound on restoration; certified restoration sequence |

Against these two works and the others in Table 5, seven interface requirements and conclusions are derived in this survey from the synthesis and are not stated in any of the reviewed sources. First, a semantic message is actionable only when it carries physical scope, validity horizon, confidence, provenance and source authority and the evidence must remain valid throughout the authorized execution interval, which is the requirement of contract I3 in Section 3.5.1. Second, an authorization that reaches the execution layer must expire no later than the earliest supporting evidence expiry after allowance for clock skew, which is contract I6. Third, operational common knowledge is a testable condition with a deadline and a default of abort, as stated in Section 4.7.2, rather than a property inferred from a broadcast. Fourth, authority is bounded by evidence through the coupling rule of Section 3.5 so that a method is admitted to a rung of the authority ladder by the evidence class it has reached and not by the technology it uses. Fifth, the six layers are complete with respect to the closed loop of Equations (1)–(3), which is argued in Section 3.5.1. Sixth, the certificate of a proof-carrying action has acceptance conditions that a verifier can evaluate under estimation uncertainty and model error, as stated in Section 7.1. Seventh, the reported outcomes of the examined streams are not comparable across streams, so comparability must be constructed by matched reporting, which is finding F4 of the principal findings stated in Section 2.4 and the reason that a proposed reference scenario with explicit implementation requirements is given in Section 7.6.1.

A compositional survey design is therefore supported by the comparison. The present work does not claim that semantic control replaces conventional telemetry, that LLM agents replace MPC or that distributed learning replaces protection. Instead, it identifies the interface conditions under which the strengths of each tradition can be combined. The mature control layer contributes deterministic dynamics, invariants and time-scale separation; distributed-control research contributes local authority and coordination mechanisms; semantic communication contributes task-relevant abstraction; agentic systems contribute deliberation and tool orchestration; and twin and assurance research contributes independent checking. The final row of Table 4 should consequently be read as an integration target and not as another algorithmic family. Its main distinction is that the unit of analysis is the complete decision lifecycle, including what an agent knows, how that knowledge was obtained, which authority it possesses, how a proposal is verified and how failures are detected and reversed.

### *2.4. Evidence Synthesis and Comparison Protocol*

A structured narrative synthesis was used because the reviewed studies address incompatible systems, data, control horizons and outcome measures. Sources were grouped by the technical question addressed, namely physical control, distributed intelligence, semantic communication, agentic reasoning, assurance or deployment. A statistical meta-analysis was not attempted because a common intervention, comparator and outcome definition were unavailable across these groups.

Each representative study was examined through the same evidence template: problem, method, control horizon, exchanged information, validation environment, inclusion of hardware-in-the-loop (HIL) testing, reported outcome, principal limitation and architectural relevance. Four power-specific comparison axes were applied: the physical horizon and affected electrical variable, the communication and common-knowledge assumptions, the location and authority of the AI component and the evidence level from conceptual analysis to field deployment. The use of preprints follows the three rules stated in Section 2.1.

The four comparison axes are made operational by a single ordinal scale that is applied to every examined study and to every comparison drawn in this survey. Five dimensions are distinguished. The horizon class (T) records the physical time scale at which the method

acts. The authority class (A) records the strongest action that the method is permitted to produce. The evidence class (E) records the strength of the validation environment. The scale class (S) records the size of the system on which the method was evaluated. The communication class (C) records what the method assumes about the channel and about the meaning of the exchanged message. The levels of each dimension are defined in Table 6. They are applied consistently in the cross-family comparison of Section 4.7, the evaluation requirements of Section 7.6 and the reproducibility and evidence-level requirements of Section 7.6.3. The same definitions are used when the maturity of a research stream is discussed.

**Table 6.** Unified evaluation scale applied to every examined study.

| Dimension | Level | Definition |
|---|---|---|
| Horizon | T1 | Sub-second: converter, protection and primary control actions |
| | T2 | Seconds to minutes: secondary restoration, balancing and fast dispatch |
| | T3 | Minutes to hours: tertiary control, energy management, market and maintenance decisions |
| | T4 | Non-real-time: planning, training, offline analysis and post-event review |
| Authority | A1 | Advisory output: text, explanation or analysis delivered to a human |
| | A2 | Structured recommendation or sandboxed proposal that is not applied to the plant |
| | A3 | Simulated or shadow action applied only inside a model or twin |
| | A4 | Approved set point applied through a verifier, shield or feasibility projection |
| | A5 | Direct actuation of a device within a proven operating envelope |
| Evidence | E0 | Conceptual analysis, architecture or position without evaluation |
| | E1 | Offline evaluation on recorded or synthetic data |
| | E2 | Closed-loop simulation with a plant model in the loop |
| | E3 | Controller hardware in the loop or real-time co-simulation |
| | E4 | Power hardware in the loop or a limited physical pilot |
| | E5 | Monitored field deployment |
| Scale | S1 | Single device, converter or converter pair |
| | S2 | Nano-grid or a single feeder with up to about ten controllable assets |
| | S3 | Microgrid or multi-feeder system with tens of buses and assets |
| | S4 | Multi-microgrid, federation or distribution network with more than one hundred buses |
| Communication | C1 | Ideal exchange: no loss, no delay and identical meaning assumed |
| | C2 | Modeled channel: delay, loss or topology represented, meaning still assumed identical |
| | C3 | Semantically explicit: ontology version, validity horizon, authority and acknowledgement represented |
| | C4 | Degraded and adversarial: mismatch, partition or attack injected and the fallback response reported |

Two comparability rules follow from Table 6 and are applied throughout the survey. The first rule is that reported numerical outcomes of different studies may be juxtaposed with their source-specific conditions but must not be used for an unmatched performance ranking. A cost reduction obtained in an energy-management simulator, a task-accuracy figure obtained on a semantic codec and a plan-success rate obtained on a language-agent benchmark are not measurements of the same quantity, they are not obtained on the same system and they are not obtained under the same communication assumptions, so those numbers cannot, by themselves, support a cross-study ranking. The second rule is that studies are compared on the basis of their class positions and on the completeness of their disclosure and that a direct comparison of outcomes requires matched tasks, data, physical assumptions, metrics, authority, communication conditions and computation budgets; broad horizon, authority and scale classes alone are insufficient. Where no matched cell exists, which is the normal situation across the streams examined here, the comparison is expressed as complementarity: what each stream establishes, what it leaves unestablished and which layer of the architecture it can therefore occupy. This is the reason that the cross-family comparison of Section 4.7 is organized by defensible claim rather than by performance and it is also the reason that the evaluation requirements of Section 7.6 specify

a common reporting set: comparability across these streams has to be constructed by future evaluations and cannot be recovered retrospectively from the published results.

The principal findings that this survey draws from the evidence are stated in Table 7, together with the basis of each finding and the part of the paper in which it is developed. The table is also intended as a reading path: the six findings carry the argument of the paper, whereas the per-cluster examination tables and the reporting checklists are the supporting evidence and the instruments through which the findings are made actionable.

**Table 7.** Principal findings of the survey, the evidence on which each rests and the part of the paper in which it is developed.

| No. | Finding | Basis | Developed in |
|---|---|---|---|
| F1 | Action authority must be allocated by layer and by evidence and it cannot be inferred from the capability of a method. | A method that plans well is not thereby qualified to actuate; the authority and evidence classes of the examined families diverge systematically. | Sections 3.5 and 4.7 |
| F2 | Semantic sufficiency and operational common knowledge are control variables and not communication conveniences. | A message without scope, validity horizon, confidence and source authority cannot be mapped to a safe action and a broadcast does not establish common knowledge. | Sections 3.2 and 4.7 |
| F3 | No examined research stream supplies all six layers, so the cognitive mini-grid is a compositional program rather than a new algorithm family. | Every stream examined in Section 4 is strong in one layer and silent in at least two others. | Sections 2.3 and 4.7 |
| F4 | Cross-study performance numbers in this field are not comparable and comparability must be constructed by matched reporting. | Studies differ in horizon, authority, scale and communication assumptions, which are the quantities that determine the reported outcome. | Sections 2.4 and 7.6 |
| F5 | Evidence in the semantic, agentic and federated streams is concentrated at levels E0 to E2, so deployment readiness is not established. | Hardware-in-the-loop, pilot and field evidence is present in the control and interface literature and largely absent in the frontier streams. | Sections 4.7 and 7.6 |
| F6 | Proof-carrying actions checked by an independent governor are the pattern that makes agentic autonomy admissible in a safety-critical mini-grid. | Assurance requires evidence that is checkable independently of the component that produced the proposal. | Sections 7.1 and 6.5 |

### *2.5. Analytical Model*

The research questions are made precise by a compact analytical model. A cognitive mini-grid can be represented as a partially observed networked dynamical system. The physical state, verified control input, disturbance and observation of agent $i$ are denoted by $x_t$, $u_t$, $w_t$ and $o_t^i$, respectively:

$$x_{t+1} = f(x_t, u_t, w_t), \tag{1}$$

$$o_t^i = h_i(x_t) + v_t^i. \tag{2}$$

In Equations (1) and (2), the state-transition function is denoted by $f$, the observation function of agent $i$ is denoted by $h_i$ and observation noise is denoted by $v_t^i$. A semantic message ($z_t^i = E_{\theta_i}(o_t^i, g_t, c_t^i)$) is generated by encoder $E_{\theta_i}$ with parameter $\theta_i$, conditioned on task $g_t$ and local context ($c_t^i$). The belief ($b_t^j$) of agent $j$ is updated from local observations and received semantic messages. An action ($a_t$) and evidence package ($e_t$) are proposed by a high-level planner, while the proposal is mapped to an executable input by a safety governor:

$$u_t = \mathcal{G}(a_t, e_t, \hat{x}_t, \Sigma_t, \mathcal{C}). \tag{3}$$

In Equation (3), the physical, cyber, contractual and organizational constraints are contained in $\mathcal{C}$ and the governor receives the state estimate ($\hat{x}_t$) and its covariance ($\Sigma_t$) from the estimator of Layer 2 rather than the physical state ($x_t$), which is not observable in the partially observed system of Equations (1) and (2); the check performed by $\mathcal{G}$ is therefore a robust check over the set of states consistent with the estimate, as specified in Section 7.1. If evidence is incomplete, stale or inconsistent, the proposal is rejected by $\mathcal{G}$ or

a fallback controller is invoked. Semantic compression, belief formation and assurance are therefore included in the closed loop rather than treated only as communication services. The formulation builds on established results rather than replacing them. The stabilizability of a plant under a limited information rate was characterized by Tatikonda and Mitter [26] and surveyed in the data-rate literature by Nair et al. [70], the trade-off between communication rate and control cost was quantified by Kostina and Hassibi [71] and stochastic networked control under information constraints was treated comprehensively by Yuksel and Basar [72]. The governor ($\mathcal{G}$) plays the role that constraint-respecting exploration plays in the safe-RL taxonomy of Garcia and Fernandez [64], while the requirement that a proposal carry checkable evidence follows the proof-carrying-code principle of Necula [73].

## 3. Concepts and Unified Cognitive Mini-Grid Architecture

In this section, the conceptual basis of the cognitive power mini-grid is developed. The progression from hierarchical control to distributed cognition is traced; the conditions that distinguish a cognitive mini-grid from an integrated one are stated; semantic control and operationally sufficient common knowledge are defined; and agent roles, tool use, memory, distributed learning and digital-twin intelligence are examined. These capabilities are then organized into a unified six-layer reference architecture in which the information flows, control time scales, authority limits, validation paths and prohibited behaviors of every layer are made explicit. The contract that couples each adjacent pair of layers is specified; the completeness of the layer set is argued from the analytical model; the correspondence between the layers, the classical control hierarchy and the normative microgrid documents is established; the relationship between the classification schemes used throughout the survey is stated; and a fault-restoration sequence is traced across all six layers as an end-to-end example. A clear separation is thereby maintained between deterministic plant control and slower artificial-intelligence-assisted deliberation.

### *3.1. From Hierarchical Control to Distributed Cognition*

Hierarchical microgrid control separates converter-level dynamics, secondary restoration and tertiary economic coordination [4]; the decentralized and hierarchical control architectures of intelligent microgrids and their power-quality, storage and AC/DC extensions were reviewed in a two-part treatment by Guerrero et al. [74,75]. This separation remains essential because voltage and frequency control cannot wait for a general-purpose reasoning model. Distributed control and MAS approaches improve resilience by allowing local controllers to cooperate using neighbor information, consensus or distributed optimization. Distributed averaging restores frequency and voltage in islanded microgrids without a central secondary controller [76], distributed control can attain economic optimality without the strict primary–secondary–tertiary hierarchy [77] and corresponding distributed control and optimization results hold for DC microgrids [78]. Cooperative multi-agent control of small-scale power networks was formulated with graph-theoretic tools by Bidram et al. [79], while the multi-agent microgrid controller of Dimeas and Hatziargyriou [11] and the agent-based DER management of Nunna and Doolla [12] demonstrated the organizational side of the same continuum.

The physical foundation of this continuum is equally well documented. Converter control in AC microgrids [80], small-signal modeling and testing of inverter-based islanded operation [81] and microgrid modeling from fundamental physics to phasor and voltage-source representations [82] describe the plant. Voltage and frequency-droop control of parallel inverters [83], adaptive droop for grid-connected and islanded operation [84], the synchronization and power-sharing conditions of droop-controlled inverters [85] and the review of power-sharing strategies for islanded AC microgrids [86] describe primary

control. Hybrid AC/DC operation is covered by the power-control and management scheme of Eghtedarpour and Farjah [87], the autonomous operation of AC and DC subgrids is studied by Loh et al. [88] and a modeling and control overview is provided by Mahmoud et al. [89]. Predictive control contributes the planning layer: the theory and computation of MPC are established in the textbook of Rawlings et al. [90], distributed and hierarchical MPC architectures were reviewed by Scattolini [91], MPC was validated on an experimental microgrid by Parisio et al. [92] and a two-layer stochastic MPC scheme addressed forecast uncertainty [93].

Distributed cognition extends this structure in three ways. First, agents maintain explicit beliefs about resources, risks, goals and the knowledge of peers. Second, messages are generated according to task significance and semantic urgency, not only sampling rates. Third, high-level decisions can be composed from specialized agents for forecasting, market negotiation, restoration, maintenance and validation, while final execution remains under local safety governors.

The transition is not from centralized to fully decentralized control. Practical systems require dynamic mixtures of local autonomy, leader-based coordination, cloud support and human authority. Holonic structures are particularly relevant because a device can act as an autonomous unit, participate in a mini-grid holon and contribute to a federation of mini-grids; semantic and holonic multi-agent management of DERs was proposed for this purpose by Howell et al. [68]. The agent abstractions on which such structures rest are those of the multi-agent textbooks of Wooldridge [94] and Shoham and Leyton-Brown [95], while the coordination of agents over a communication graph is analyzed by consensus theory [14] and by graph-theoretic methods for multi-agent networks [96].

The underlying coordination mechanisms are supplied by several lines of work. Agent-based software engineering provides the design discipline [97]. Consensus seeking under dynamically changing interaction topologies [98], distributed subgradient methods for multi-agent optimization [99] and distributed optimization through the alternating direction method of multipliers (ADMM) [100] provide the algorithms, which are surveyed more broadly in the distributed-optimization literature [101]. Microgrid demonstrations show their practical roles in dispatch and restoration. A MAS for real-time microgrid operation was implemented in a real-time digital simulator [102] and distributed smart-grid agents were shown to isolate and supply critical loads after an upstream outage [103]. Cooperative multi-agent control coordinated heterogeneous storage in a DC microgrid [104] and asynchronous consensus ADMM solved microgrid energy management under communication imperfections [105]. The case for distributed intelligence as an essential element of the microgrid has also been made from an industry perspective [106].

### Defining Conditions of a Cognitive Mini-Grid and the Distinction from Integration

A mini-grid that combines multi-agent control, semantic communication, digital twins and language-model agents is an integrated mini-grid; it is not thereby a cognitive one. The distinction is stated here as four conditions, each of which can fail in an integrated system and each of which is a property of the closed loop rather than of any component. First, the epistemic state of the agents is a control variable: a coordinated action is admitted only when operational common knowledge of the relevant state has been established by acknowledgment or by a trusted replicated record and never on the grounds that a message was broadcast. Second, semantic sufficiency is a contract: a message is actionable only when it carries its physical scope, validity horizon, confidence, provenance and source authority; a message without them is a datum that may inform a belief but cannot license an action. Third, authority is separated from capability by proof-carrying action: a deliberative component may produce proposals of any sophistication, but a proposal reaches the plant

only as an evidence package checked by a governor that is independent of the proposer and bounded by the evidence class of the function. Fourth, learning is closed-loop and governed: outcomes update forecasts, policies and twins under a registry that records data scope, version and validation conditions and no update alters a safety policy without passing the same governor. The eight capabilities enumerated in Section 1 are the capabilities that such a system possesses; the four conditions are what make their composition cognitive rather than merely integrated.

The distinction is made operational in Table 8, which contrasts the typical properties of an integrated smart-grid architecture with the properties required here and names the layer in which each is realized. The table also supplies a test: a system that composes all of the enabling technologies but that conditions a coordinated action on a broadcast alone, acts on a message that carries no validity horizon, allows a planner to actuate without an independent check, or updates a policy outside the governor satisfies the description of integration and fails the description of cognition. None of the four conditions requires a new algorithm and none is claimed as an invention of this survey. The claimed contribution is their joint statement as necessary conditions, the assignment of each condition to a layer of the reference architecture and the proof-carrying action pattern through which the third condition is enforced.

**Table 8.** Properties that distinguish a cognitive mini-grid from an integrated smart-grid architecture and the layer in which each is realized.

| Property | Integrated Smart-Grid Architecture, Typical | Cognitive Mini-Grid, Required | Realized in |
|---|---|---|---|
| Basis of a coordinated action | Receipt of a broadcast or of a shared signal | Established operational common knowledge with acknowledgement, validity interval and abort rule | Layer 4 |
| Status of a message | Data point with an assumed common meaning | Semantic object carrying scope, validity horizon, confidence, provenance and authority; not actionable without them | Layer 3 |
| Relation between planning and execution | Planner output applied as a set point, at most after a feasibility check by the same component | Proposal with an evidence package, admitted only by an independent governor and bounded by the evidence class of the function | Layers 5 and 6 |
| Treatment of model output | Forecast or inference consumed as if it were measured | Confidence and provenance carried with the value and reflected in the permitted authority | Layers 3 and 6 |
| Learning and adaptation | Offline retraining and redeployment | Governed continual learning with registry and versioning; no unreviewed change to a safety policy | Layers 4 and 6 |
| Failure of an intelligent component | Manual fallback or undefined behavior | Degradation to the highest satisfied contract while protection and local control continue | Layers 2 and 6 |
| Unit of certification | The controller function | The controller function, together with the evidence required before an externally originated action is admitted | Layer 6 |

### *3.2. Semantic Control and Common Knowledge*

Distributed cognition depends on what the agents exchange. Semantic communication seeks to preserve task-relevant meaning rather than every source bit. In mini-grids, this suggests an encoder optimized against control loss, risk or restoration performance. A generic objective is

$$\min_{E,D,\pi} \mathbb{E}[L_{\text{task}}(x_t, \tilde{x}_t, u_t) + \lambda R(z_t) + \mu L_{\text{risk}}]. \quad (4)$$

In Equation (4), communication cost is measured by $R(z_t)$, the decoder is denoted by $D$, the control policy is denoted by $\pi$ and the task and risk losses are denoted by $L_{\text{task}}$ and $L_{\text{risk}}$. The decoded state at the receiving agent is $\tilde{x}_t = D(z_t)$, which is distinct from the estimate ($\hat{x}_t$) of Equation (3) produced by the estimator of Layer 2 and $\pi$ generates a candidate action whose verified realization is $u_t$ through Equation (3). The non-negative weighting coefficients are denoted by $\lambda$ and $\mu$. The message alphabet may contain continuous latent vectors, discrete intent codes, ontology-linked events or natural-language statements. Shannon's theory deliberately excluded meaning from the engineering problem [107]; Weaver subsequently distinguished the technical, semantic and effectiveness levels of communication [108] and a formal theory of semantic information was outlined by Carnap and Bar-Hillel [109]. The modern semantic-communication program builds on these foundations. Bao et al. proposed a model-theoretic framework for semantic communication [19] and Strinati and Barbarossa positioned semantic and goal-oriented communication beyond Shannon for 6G [20]. Lan et al. clarified what conveying meaning entails for machine intelligence [21], Gunduz et al. unified context, semantics and task-oriented communication [22] and Xie et al. demonstrated a learned end-to-end semantic transmission system [23]. Unlike ordinary compression, a semantic message in a mini-grid must additionally specify the context, confidence, source authority and validity horizon; this requirement is a proposal of the present survey rather than a result of the cited works. Related mechanisms are available from neighboring fields. Deep joint source-channel coding transmits images directly over noisy channels [110] and exploits channel feedback [111]. Edge inference has been formulated as wireless image retrieval [112] and as task-oriented communication for multi-device cooperative inference [113]. Ultra-reliable low-latency communication has been analyzed from a communication-theoretic network-slicing perspective [114] and the prospects of semantic communication for 6G have been assessed [57]. These mechanisms supply useful building blocks, but mini-grid semantics additionally require physical validity, temporal scope and authorization.

Common knowledge is stronger than data sharing. A fact is common knowledge when every relevant agent knows it, knows that others know it and so on. Exact common knowledge is unattainable in a distributed system whose communication is not reliable and instantaneous, as shown by Halpern and Moses [115], who also introduced the weaker, attainable variants on which practical protocols rest; operationally sufficient approximations can be created through authenticated broadcast, acknowledgment, synchronized observations and public commitments. A restoration plan should begin only when the participating switch, storage and load agents have a mutually consistent view of the fault boundary and authority to act. The testable form of this condition, with its protocol and failure assumptions, is given in Section 4.7.2. MACKRL demonstrates the coordination value of conditioning hierarchical policies on common knowledge [15], although direct mini-grid implementations remain largely unexplored.

This leads to an epistemic safety principle: an action may be physically feasible under one agent's state estimate yet unsafe when peer knowledge is inconsistent. Cognitive control must therefore evaluate both physical feasibility and the epistemic preconditions for coordinated execution.

### *3.3. Agentic Roles, Tools and Memory*

The reasoning that evaluates those preconditions is increasingly delegated to agentic components. Agentic AI transforms a language or multimodal model into a component through which goals can be pursued with planning, memory and tool use. A single unconstrained "super-agent" should not be employed in a cognitive mini-grid. Roles and privileges should be separated. Goals are decomposed by a planner. Measurements and

documents are retrieved by an analyst and uncertainty-aware scenarios are provided by a forecasting agent. Negotiation is performed by a market agent and counterfactuals are evaluated by a digital-twin agent. Constraints are checked by a validator and bounded commands are issued by an executor. Tool-use methods show how reasoning can be grounded in external functions and knowledge bases: ReAct interleaves reasoning traces with tool actions [31], RAG conditions generation on retrieved documents [30] and Toolformer-style interfaces let a model learn when and how to call external tools [32]. Reflexion adds verbal self-reflection across episodes [116], while AutoGen [44] and MetaGPT [45] organize role-specialized agents into conversational or meta-programmed workflows. The broader technical context is established by few-shot-capable foundation language models [117], open-weight model families [59], generative-agent architectures that maintain memory and plans [118] and surveys of the opportunities and risks of LLM-based agents [46].

Grid-Agent [33] and X-GridAgent [34] illustrate the emerging pattern of combining natural-language reasoning with domain-specific numerical tools and hierarchical agent layers, a pattern also identified in the smart-grid review of Kiasari and Aly [35]. Their value is less in replacing optimal power flow than in orchestrating analyses, interpreting goals and producing transparent workflows. Their limitations, as assessed in this survey, include preprint status, restricted test environments, uncertain latency and the absence of established certification pathways.

Memory must also be engineered. Short-term memory stores the current incident, tool outputs and approvals. Long-term memory stores asset histories, policies and prior cases. Learned memory should be separated from authoritative records so that an agent cannot overwrite protection settings or contractual rules through conversational interaction. Retrieval results require provenance and versioning.

### *3.4. Learning and Digital-Twin Intelligence*

Agents also adapt; adaptation raises the question of where learning takes place and how its outputs are validated. Learning occurs at multiple scopes. Device-level models estimate the state of charge, degradation or local disturbances. Mini-grid models forecast demand and renewable output. At the energy-management level, RL has been applied to distributed economic management [119], to dynamic energy management with recurrent value approximation [120], to real-time scheduling with deep RL [121], to multi-microgrid coordination [122] and to safety-constrained energy management in distribution networks [123]. The RL literature for microgrid control was surveyed by Barbalho et al. [50]. MARL coordinates multiple resources. FL enables several mini-grids to share model updates without transferring raw consumer data, following the federated averaging algorithm [16] and subject to the statistical and systems challenges reviewed by Li et al. [124]. Graph neural networks encode electrical topology [18] and may support transferable foundation models, a direction explored for grid operations [52] and for federated foundation models shared across organizations [125]. The algorithmic baseline spans the value-based, actor–critic, deterministic-policy and entropy-regularized families of single-agent RL, namely the textbook foundations [126], deep Q-networks [127], deep deterministic policy gradient [128], the twin-delayed variant that addresses function-approximation error [129], proximal policy optimization [49] and soft actor–critic [130]. Cooperative multi-agent learning adds centralized-critic training [131], counterfactual credit assignment [132] and value factorization [133,134], as surveyed critically by Hernandez-Leal et al. [135]. Federated machine learning, as a concept and its application classes were formalized by Yang et al. [136]. These families remain baselines rather than deployment guarantees.

A digital twin is the principal environment for safe deliberation. The concept originated in the aerospace paradigm of a high-fidelity virtual counterpart of a physical

vehicle [137]; its industrial state of the art was reviewed by Tao et al. [61] and its enabling technologies and open challenges were surveyed by Fuller et al. [138]. Kritzinger et al. distinguished the digital model, the digital shadow and the digital twin by the degree of automated data exchange [139]. Jones et al. characterized the twin through a systematic literature review [140] and Rasheed et al. examined its values, challenges and enablers from a modeling perspective [141]. In the power domain, lessons from laboratory-scale live twins of DC microgrids have been reported [142]. Such a twin integrates physical models, learned surrogates, asset histories and communication conditions to evaluate proposed actions before execution. A causal twin goes beyond correlation by representing intervention effects and supporting counterfactual questions in the sense of the structural causal models of Pearl [143], the causal-inference foundations of Peters et al. [144] and causal representation learning [145]. For example, it should distinguish between a voltage decline caused by renewable variability and one caused by a malicious measurement because the appropriate intervention differs. Hybrid twins combine interpretable physical models with data-driven residuals and uncertainty quantification, following the hybrid analysis and modeling advocated by Rasheed et al. [141]. The digital-twin concept and its smart-grid specialization further motivate lifecycle models that connect design, testing, operation and maintenance [60,146].

Federated foundation models are a longer-term direction rather than a present deployment solution. The foundation-model paradigm itself, with its opportunities and risks, was characterized by Bommasani et al. [29]. Such models could learn representations across time series, topology, images and maintenance logs but require energy-aware training, data governance, continual adaptation and uncertainty layers. Recent work has begun to frame the opportunity for grid operations [52] and for federated foundation models shared across grid organizations [125] and to add a conformal-prediction trustworthiness layer to a power-system foundation model [54].

### *3.5. Unified Six-Layer Architecture*

These capabilities are brought together in a single reference architecture. The cognitive mini-grid is organized into the six layers shown in Figure 2. The physical layer contains energy assets. The sensing and edge-reflex layer contains measurement, protection and fast control. The semantic and common-knowledge layer transforms data into task-relevant beliefs and commitments. The distributed-intelligence layer coordinates specialized agents and learning. The agentic-deliberation and digital-twin layer performs slower planning and counterfactual evaluation. The assurance, governance and cyber-resilience layer enforces safety, cybersecurity and governance across all lower layers.

The layer comparison in Table 9 makes the authority boundary of the architecture explicit; its last column lists representative papers for each layer. Tables 4 and 9 serve different purposes and share no unit of analysis: Table 4 is keyed by review tradition and records what each body of literature covers and omits, whereas Table 9 is keyed by layer and records what each part of the proposed architecture is responsible for and forbidden from doing. The physical-energy and edge-reflex layers are closest to the plant and therefore operate with the strongest deterministic constraints and the shortest deadlines. Their prohibited behaviors are deliberately strict: they must not wait for cloud reasoning, accept unauthenticated commands or delegate stability-critical responses to an LLM. The semantic and distributed-intelligence layers have more representational freedom, but their outputs remain descriptions, commitments, forecasts or candidate coordination decisions rather than direct unrestricted actuation. The agentic/twin layer has the broadest deliberative capability because it can retrieve information, call tools and compare scenarios, yet it is also the layer most explicitly prohibited from direct actuation. The assurance and governance

layer is placed above the full stack not because it operates at the fastest rate but because it defines the policies, authorization and evidence requirements that every downward action must satisfy.

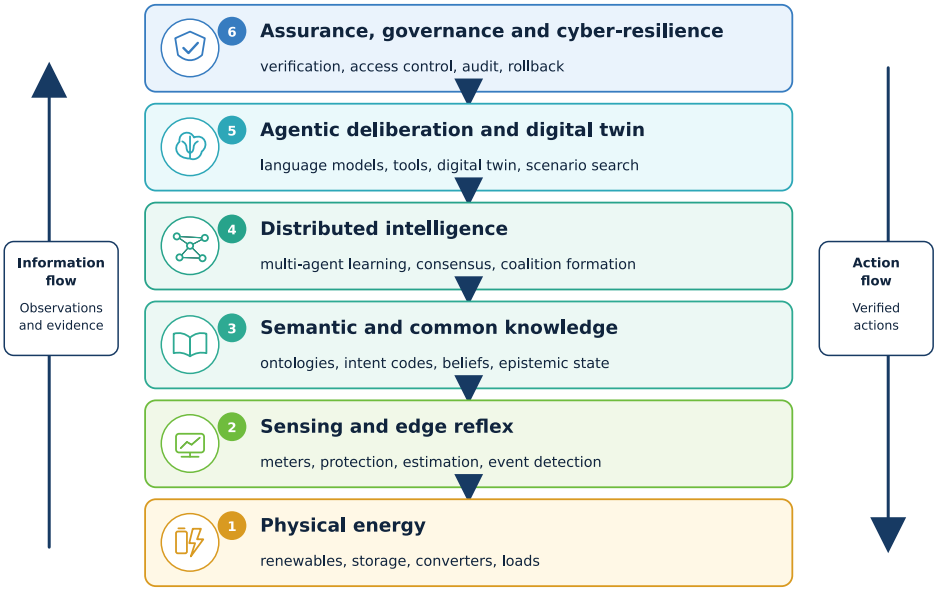


**Figure 2.** Unified cognitive power mini-grid architecture, with observations and evidence moving upward and verified actions moving downward.

**Table 9.** Functional responsibilities and prohibited behaviors by architecture layer.

| Layer | Primary Functions | Representative Methods | Must Not Do | Representative Papers |
|---|---|---|---|---|
| Layer 1, Physical energy | Generate, store, convert and consume power | Inverters, batteries, switches, relays and flexible loads | Accept unauthenticated or unconstrained high-level commands | [1,2,80,147,148] |
| Layer 2, Sensing and edge reflex | Protection, state estimation and fast control | Droop control, MPC, local observers and event-triggered control | Depend on cloud or LLM latency for stability-critical response | [4,5,83,92,93] |
| Layer 3, Semantic and common knowledge | Encode intent, urgency, confidence and shared commitments | Ontologies, semantic codes, acknowledgments and epistemic logic | Remove information needed to verify safety or source authority | [15,19,20,22,27,68] |
| Layer 4, Distributed intelligence | Coordinate local agents and learn shared policies | MAS, consensus, ADMM, MARL, FL and graph learning | Assume identical observations, models or communication quality | [9,11,16,18,105,119,149] |
| Layer 5, Agentic deliberation and digital twin | Plan, retrieve, call tools, evaluate scenarios and explain | Language-model agents, RAG, digital twins and causal models | Directly actuate assets without independent validation | [30,31,33,34,60,142,143] |
| Layer 6, Assurance, governance and cyber-resilience | Verify constraints, authorize, audit and roll back | Neuro-symbolic rules, formal methods, runtime assurance and access control | Allow self-modification of safety policies by an untrusted agent | [55,62,73,150–152] |

Along the architecture stack, information richness and action authority move in opposite directions. Upward flows may become progressively more abstract: raw samples are transformed into events, semantic states, beliefs, forecasts and candidate plans. Downward flows must become progressively more concrete and more constrained: a high-level objective is translated into a validated schedule; an authorized set point; and, finally, a bounded device command. This asymmetry prevents a common design error in which a flexible semantic or language interface is granted the same authority as a certified controller. It also creates clear failure-containment rules. When an ontology is mismatched, a model is unavailable or a digital twin loses synchronization, the affected upper layer can be disabled

while protection and local control continue. Table 9 therefore serves as both a functional decomposition and a safety allocation: every new cognitive function must be assigned a layer, a permissible time horizon, an input-evidence contract and an explicit list of actions that it is not allowed to perform.

The four capability-maturity levels, denoted as M1 to M4 in the text and labeled Stage 1 to Stage 4 in Figure 3, classify increasing cognitive capability. They are independent of the authority and evidence classes defined in Table 6. Functional breadth is described by maturity, permitted actuation by authority and validation strength by evidence level. No one-to-one mapping is assumed, but the schemes are not unrelated and their relationship is stated in Table 10 so that they can be kept apart. Four classification schemes are used in this survey and they classify four different objects. The maturity level classifies a capability by its breadth along the six dimensions of Figure 3. The authority class classifies an output by the strongest effect it may have on the plant. The evidence class classifies a study or a deployed function by the strength of its validation. The roadmap stage of Section 7.6.4 classifies the research and deployment program by the order in which its objectives are pursued. Two coupling rules connect them. The first rule couples authority to evidence for every function of Layers 3 to 6: an approved set-point (A4) requires at least controller hardware-in-the-loop evidence (E3) before a pilot; direct actuation (A5) requires power hardware-in-the-loop or pilot evidence (E4); and advisory, recommending and simulated outputs (A1 to A3) are admissible for any evidence class. The functions of Layers 1 and 2 hold their authority through the certification path of the standards mapped in Section 3.5.2. The second rule couples the roadmap stage to the available evidence and to the authority admissible at that stage: Stage 1 works at E0 to E1 and admits, at most, A2; Stage 2 works at E2 and admits, at most, A3; Stage 3 works at E3 to E4 and admits, at most, A4; Stage 4 requires E4 to E5 with certification and admits A4 to A5; and Stage 5 requires E5 across federated sites and admits A5 across them. Maturity is coupled to neither rule: a capability of any maturity level may be studied at any stage and a Stage 1 project may prototype an M4 capability at E1 with A1 authority. The horizon, scale and communication classes of Table 6 are descriptive dimensions of a study and not classifications of progress, so they have no counterpart in the table.

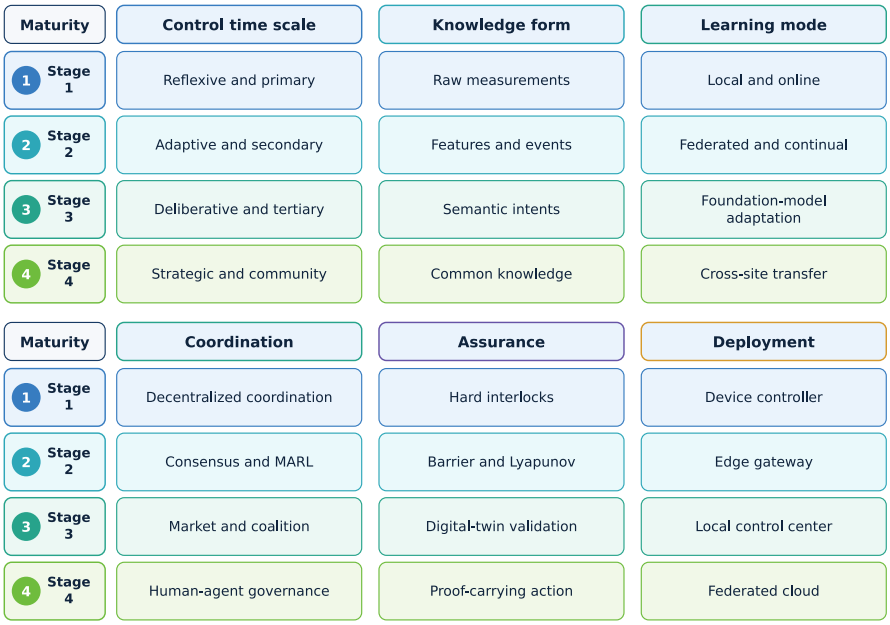


**Figure 3.** Taxonomy linking control timescale, knowledge form, learning mode, coordination, assurance and deployment location across the four capability-maturity levels (M1 to M4), which are labeled Stage 1 to Stage 4 in the figure and are distinct from the roadmap stages of Section 7.6.4.

**Table 10.** Relationship between the four classification schemes used in this survey.

| Scheme | What It Classifies | Levels | Relation to the Other Schemes |
|---|---|---|---|
| Capability maturity (Figure 3) | A capability, or a mini-grid, by its breadth along control time scale, knowledge form, learning mode, coordination, assurance and deployment location | M1 to M4 | Independent of the other three: it describes what can be represented and done, not what may be actuated or how well it is validated |
| Authority class (Table 6) | An output or action, by the strongest effect it is permitted to have on the plant | A1 to A5 | Bounded above by the evidence class through the first coupling rule; assigned per function and per use case, never per technology |
| Evidence class (Table 6) | A study, or a deployed function, by the strength of its validation | E0 to E5 | Prerequisite for authority; determines the roadmap stage at which a function may be pursued |
| Roadmap stage (Section 7.6.4) | The research and deployment program, by the order in which its objectives are pursued | Stage 1 to Stage 5 | Bounded by the available evidence and the authority admissible through the second coupling rule; independent of maturity |

#### 3.5.1. Inter-Layer Coupling and Interface Contracts

The layers are coupled through a small number of explicit contracts rather than through unrestricted access to one another. Each adjacent pair is connected by exactly one contract and each contract states what is passed upward, what is passed downward, at which nominal rate, within which latency budget, in which representation and what is done when the contract cannot be satisfied. The contracts are listed in Table 11. Two properties are enforced by this arrangement. The first property is that authorization and evidence contracts cannot be bypassed: a deliberation result is passed through the governor before the explicit I6 return path from Layer 6 to Layer 2 is used, and a raw sample is interpreted with its semantic metadata before it is incorporated into a belief. The second property is that the failure of an upper contract degrades the system to the highest layer that is still satisfied so that the loss of the twin, of the ontology or of the language model leaves protection and local control operating on the contracts of the two lowest layers.

The completeness of the layer set is argued from the analytical model of Section 2.5 rather than asserted and it is argued as coverage and separation rather than as a uniqueness proof. The closed loop of Equations (1)–(3) contains six objects that must be realized by some part of a deployed system: the plant dynamics ($f$); the observation map ($h_i$), together with the estimator that produces $\hat{x}_t$ and $\Sigma_t$ from it; the semantic encoder ($E_{\theta_i}$); the belief ($b_t^j$); the pair of proposed action and evidence $(a_t, e_t)$; and the governor $\mathcal{G}$. Coverage holds because each object is assigned to one layer and to one only. The plant to Layer 1, the observation map and its estimator to Layer 2, the encoder to Layer 3, the belief and its distributed formation to Layer 4, the proposal and its evidence to Layer 5 and the governor with the constraint set ($\mathcal{C}$) to Layer 6, so that no object of the loop is left unrealized and no layer is without an object. Separation holds because adjacent objects differ in the property that a contract must protect: time scale between Layers 1 and 2, meaning between Layers 2 and 3, epistemic state between Layers 3 and 4, authority between Layers 4 and 5 and trust between Layers 5 and 6. The corresponding logical contracts must be preserved, even when several layers are implemented within one component, which motivates the proposed separation. The quantities that might appear to require further layers. Namely, the task ($g_t$), the local context ($c_t^i$) and the organizational and contractual constraints contained in $\mathcal{C}$, are inputs to the loop rather than transformations within it, so markets, operators and regulators are represented as sources of $g_t$ and of $\mathcal{C}$ and not as additional layers. The argument is conditional on the model: it establishes that the six layers cover the closed loop of Equations (1)–(3) and that each boundary protects a distinct property and it does not establish that no other partition could do so.

**Table 11.** Interface contracts between adjacent layers of the architecture.

| Interface | Passed Upward | Passed Downward | Nominal Rate and Latency Budget | Behavior When the Contract Fails |
|---|---|---|---|---|
| I1, Layer 1 to Layer 2 | Sampled electrical quantities, switch and breaker states and device health | Gate signals, references and trip commands | Sub-cycle to seconds; deterministic execution required | Protection settings and local droop remain in force; the device is isolated if measurements are lost |
| I2, Layer 2 to Layer 3 | Validated estimates, events and quality flags with timestamps | Acknowledged set points and mode requests within the device envelope | Seconds; bounded queueing delay | Semantic encoding is suspended and the layer above receives no state, which downgrades every dependent action |
| I3, Layer 3 to Layer 4 | Semantic states, intents, commitments and confidence with validity horizon and ontology version | Requests for evidence, refresh or acknowledgement | Seconds to minutes; the validity horizon of the evidence must cover the authorized execution interval | Ontology mismatch is declared, the message is rejected and coordination reverts to raw telemetry at reduced authority |
| I4, Layer 4 to Layer 5 | Agreed beliefs, coordination outcomes, learned policies and forecasts | Task assignments, constraint updates and negotiation mandates | Minutes; asynchronous | Deliberation proceeds on stale beliefs only if the staleness is declared; otherwise, the plan request is refused |
| I5, Layer 5 to Layer 6 | Candidate action, assumptions, predicted consequences, validity horizon and evidence package | Verdict, required additional evidence and permitted authority level | Minutes; no deadline on the plant and a proposal whose evidence expires before a verdict is returned unexecuted | An incomplete evidence package is rejected without partial execution |
| I6, Layer 6 to Layer 2 | Execution outcome, deviation and audit record | Authorized and bounded set point with expiry and rollback condition | Seconds; the authorization expires no later than the earliest expiry of the evidence based on which it was granted, allowing for clock skew | Execution is not started after expiry; if a sequence has begun, the certified recovery rule of Section 7.1 is evaluated from the measured intermediate state and the previously verified controller resumes |

#### 3.5.2. Relation to Established Hierarchies and to Standard Microgrid Controller Functions

The architecture is a refinement of the established hierarchy and not a replacement of it. The correspondence is given in Table 12. Primary and secondary control are mapped onto Layer 2, while tertiary coordination is mapped onto Layer 4, with semantic information exchanged through Layer 3. The transition and dispatch functions of the microgrid controller specified in IEEE Std 2030.7 [153] are assigned to Layers 2 and 4, respectively. This is a functional mapping rather than a demonstrated compliant implementation. Test procedures for those controller functions are provided by IEEE Std 2030.8 [154]; compliance must be assessed for the implemented controller. Layers 1 and 2 are constrained by the interconnection requirements of IEEE Std 1547 [155]. The vocabulary of the semantic layer is drawn from the distributed-energy-resource information model of IEC 61850-7-420 [156] and the common information model of IEC 61970 and IEC 61968 [157]. The security requirements of the IEC 62351 series [151] are relevant across the stack. Security guidance is also provided by NIST Interagency Report (NISTIR) 7628 [152] and NIST Special Publication (SP) 800-82 [158].

**Table 12.** Correspondence between the proposed layers, the classical control hierarchy and the normative microgrid documents.

| Layer | Classical Hierarchy | Normative Counterpart | What the Layer Adds Beyond the Normative Counterpart |
|---|---|---|---|
| Layer 1, Physical energy | Plant and converters | IEEE Std 1547 interconnection and performance requirements [155] | Nothing is added; the layer is constrained by the standard |
| Layer 2, Sensing and edge reflex | Primary and secondary control | Transition function of IEEE Std 2030.7 and its tests in IEEE Std 2030.8 [153,154] | An expiry and a rollback condition are attached to every externally originated set point |
| Layer 3, Semantic and common knowledge | No counterpart | Information models of IEC 61850-7-420 and IEC 61970 [156,157] | Validity horizon, confidence, source authority and ontology version are made part of the message rather than of the documentation |
| Layer 4, Distributed intelligence | Tertiary coordination | Dispatch function of IEEE Std 2030.7 and the standards and practice of multi-agent power engineering [10,153] | Operational common knowledge is defined and tested before a coordinated action is admitted |
| Layer 5, Agentic deliberation and digital twin | No counterpart | No operational counterpart; the offline design activity is covered by IEEE Std 2030.9 [159] | Deliberation is admitted as a proposal generator with no execution authority |
| Layer 6, Assurance, governance and cyber-resilience | No counterpart | IEC 62351, NISTIR 7628 and NIST SP 800-82 for security; the artificial-intelligence risk-management framework and ISO/IEC 42001 for governance [55,151,152,158,160] | A machine-checkable evidence package is required before authority is granted and the check is performed independently of the proposing component |

Three additions distinguish the architecture from the normative hierarchy and they are additions rather than disagreements. First, IEEE Std 2030.7 specifies what the microgrid controller must do and how its functions are tested, but it does not specify the provenance, the validity horizon or the semantic version of the information based on which a dispatch decision is taken because that information was assumed to originate from deterministic supervisory acquisition; Layer 3 supplies that specification. Second, the normative documents contain no admissibility condition for a proposal produced by a learned or generative component, since no such component is contemplated in them; Layers 5 and 6 supply the separation between a proposal and an authorized action, together with the evidence that must accompany the proposal. Third, the standards define the functions of a single microgrid controller rather than the epistemic conditions under which several autonomous

agents may commit shared resources; Layer 4 supplies the definition of operational common knowledge and its establishment test. The two tables that treat the standards are keyed differently and are not interchangeable: Table 12 is keyed by layer and states where each normative asset already applies, whereas the profile table of Section 7.6.4 is keyed by asset and states what each asset must be extended with. Conversely, nothing in the architecture relaxes a normative requirement: the dispatch and transition functions retain their deterministic specification, the interconnection requirements are unaltered and the added layers can only reduce (never extend) the authority that the standard already grants to the microgrid controller.

#### 3.5.3. End-to-End Illustrative Example: Fault Isolation and Restoration

The contracts of Table 11 and the coupling rules of Table 10 are illustrated by one operating sequence: a permanent fault in a feeder zone of a grid-connected community mini-grid during a period of high photovoltaic infeed, followed by the restoration of the disconnected loads. The sequence is chosen because it crosses every layer and both directions of information flow and because it depends on the physical resources whose planning is discussed in Section 5.2. The steps are listed in Table 13. The quantities in the table are illustrative values chosen to make the messages concrete and the times are the latency budgets of the corresponding contracts rather than measurements.

Three points of engineering clarity follow from the sequence. First, capability and authority are visibly separated: the most capable component, the planner of step 6, never touches the plant and the only component that does, the protection and control functions of steps 2 and 8, is the least deliberative. Second, the physical resources determine what any layer can achieve: the plan of step 6 exists only because storage B2 holds 240 kWh, a sectionalizing switch (SW7) was installed and a mobile storage unit was pre-positioned and the revision of step 7 is possible only because the mobile unit can be connected first, which are planning decisions of the kind discussed in Section 5.2 and the cognitive layers can allocate those resources but cannot create them. Third, the evidence and authority classes fix what may be automated at each stage of the roadmap: the sequence could be considered for a Stage 3 pilot at evidence class E3 with authority A4 at step 8 only after every acceptance condition has been verified, whereas in a Stage 2 prototype, the same sequence stops at step 7 with a simulated action at A3 and in a Stage 1 study, it produces the advisory output of step 6 only.

**Table 13.** Fault isolation and restoration traced across the six layers: information flow, agent decisions and certified execution.

| Step | Layer and Contract | Information Carried or Decision Taken | Authority and Time Budget |
|---|---|---|---|
| 1 | Layer 1 | Fault current in feeder zone Z3; the recloser trips and the sectionalizing switch (SW3) opens | Physical event; sub-cycle |
| 2 | Layer 2, contract I1 | Protection isolates Z3 with the recloser and SW3; droop and local controllers hold voltage and frequency on the remaining sections; state estimation flags the measurements that were lost with the section | A5 by the certified protection and control functions; sub-cycle to seconds |
| 3 | Layer 2 to Layer 3, contract I2 | Validated event: section identifier, isolated loads, remaining generation and the state of charge of storage, all timestamped and carrying quality flags | Information only; seconds |
| 4 | Layer 3, contract I3 | Semantic event and intent objects: zone Z3 isolated; 180 kW of critical load unserved; stored energy of 240 kWh at storage B2 at 60 percent confidence; ontology version 4; validity horizon of 5 min; source authority, substation agent | Information only; the validity horizon covers the authorized execution interval |

**Table 13.** *Cont.*

| Step | Layer and Contract | Information Carried or Decision Taken | Authority and Time Budget |
|---|---|---|---|
| 5 | Layer 4, contracts I3 and I4 | The agents of the storage units, of the sectionalizing switches and of the mobile storage unit exchange beliefs; operational common knowledge of the isolated section, of the available reserve and of the switch states is established by acknowledgment; the restoration task is assigned | Coordination outcome, A2; seconds to minutes |
| 6 | Layer 5, contracts I4 and I5 | Restoration plan proposed: reconfigure through switch SW7, dispatch storage B2 at 150 kW, route the mobile storage unit to bus 14 and shed the deferrable loads; the digital twin evaluates the plan against voltage, thermal and reserve constraints; the evidence package is assembled with assumptions, predicted consequences, validity horizon and certificate | Proposal, A2 to A3; minutes; no deadline on the plant |
| 7 | Layer 6, contract I5 | The governor checks the certificate against the constraint set: switch interlocks, protection coordination for the new topology, reserve policy, authorization of the proposing agent and freshness of the evidence; the first submission is returned with the failed conditions and their margins, worked through in Section 7.1; the revised illustrative submission, which connects the mobile unit first, reduces the dispatch and shortens the expiry, satisfies the stated scalar margins but requires the remaining checks before authority A4 with a 4-min expiry can be granted | Verdict; minutes |
| 8 | Layer 6 to Layer 2, contract I6 | If all checks for the revised plan pass, its authorized set points and switching sequence are delivered with expiry and certified recovery conditions; Layer 2 executes within its envelope and continuing dispatch is reauthorized before expiry; outcomes and deviations are recorded | A4, bounded set points; seconds; expiry within the validity horizon of the evidence |
| 9 | Layers 4 and 5 | The outcome is recorded; the forecasts of reserve consumption and of restoration duration are updated under the model registry; no safety policy is altered | Governed learning; non-real-time |
| 10 | All layers, failure cases | If the acknowledgment of the mobile storage agent is missing at step 5, common knowledge is not established and the plan is reduced to the resources whose state is known; if the ontology versions differ at step 4, the message is rejected and coordination reverts to raw telemetry at authority A3; if the twin loses synchronization at step 6, the proposal is refused and protection and local control continue | Degradation to the highest satisfied contract |

## 4. In-Depth Examination of the Literature

In this section, the principal research streams from which cognitive power mini-grids are formed are examined in depth. Microgrid control foundations, learning-based coordination, semantic communication, LLM-based agents, digital twins with neuro-symbolic assurance and federated foundation models and the distributed artificial intelligence (DAI) and extended Belief–Desire–Intention (BDIx) agent lineage with flexible power-interface control are examined in turn. Each representative study is examined through a common evidence template covering the operational problem, technical method, control hierarchy, exchanged information, validation environment, reported outcome, principal limitation and architectural relevance. Hardware-in-the-loop testing is distinguished from simulation and conceptual analysis. The research families are then compared by control horizon and authority, by communication and knowledge assumptions and by safety and cybersecurity evidence so that direct comparison is enabled without ranking studies that address different technical problems. The methodological limitations, the experimental limitations and the

boundaries within which the conclusions of each family remain valid are appraised before the resulting evidence gap is stated.

### *4.1. Microgrid Control, Distributed Optimization and Multi-Agent Foundations*

Following the evidence template established in Section 2.4, the examination starts with the control foundations on which every later family builds. Lasseter's early microgrid formulation established the essential system concept: a bounded cluster of loads and microsources should appear as a controllable entity to the wider power system while maintaining local service in islanded operation [1]. The approach relies on peer-to-peer local controllers with plug-and-play capability so that each microsource regulates its own power, voltage and frequency contribution without a dedicated central computer or a fast communication network. The methodology is conceptual and design-oriented; control requirements, example configurations and separation behavior are analyzed rather than a single quantitative benchmark. The reported outcome is architectural feasibility: rapid disconnection under upstream disturbances, autonomous islanded service and coordinated reconnection were shown to be achievable with local intelligence alone. The paper's enduring novelty is therefore architectural rather than algorithmic. It clarified that local control, rapid disconnection and coordinated reconnection are first-class requirements. For the cognitive mini-grid, this remains the physical contract: semantic or agentic layers may modify objectives and set points, but they must not bypass the electrical boundary conditions that make the microgrid a coherent controllable system.

Guerrero et al. formalized hierarchical control for droop-controlled AC and DC microgrids by separating primary power sharing, secondary restoration and tertiary power-flow management [4]. The approach adapts an industrial control hierarchy to converter-dominated systems: droop and virtual-impedance loops share power without communication, a slower secondary loop removes the frequency and voltage deviations that droop introduces and a tertiary loop manages the power exchange with the external grid. The methodology combines small-signal modeling with simulation results and experimental verification on representative configurations. The reported results show that secondary action restores nominal frequency and voltage while droop-based sharing is preserved and that controlled synchronization with the main grid becomes possible; the novelty is a standardization-oriented three-level template that remains the reference decomposition for microgrid autonomy. Bidram and Davoudi subsequently systematized the hierarchical structure and the roles of centralized and distributed secondary or tertiary controllers [5]. Their methodology is a comparative survey. The principal result is a taxonomy of controller functions, communication requirements and realization options that still supplies the standard vocabulary for controller responsibilities. Olivares et al. reviewed the evolution toward advanced energy management, predictive control and communication-supported coordination through a task-force synthesis whose main outcome is a structured research agenda rather than a new algorithm [6]. Read together, these papers define the time-scale and authority boundaries that an AI survey must preserve. Their limitation relative to cognitive autonomy is not a deficiency in the original work; it is that semantic meaning, epistemic uncertainty and tool-using agents were outside their scope.

Parisio et al. demonstrated the practical value of MPC for experimental microgrid optimization by explicitly forecasting future trajectories and enforcing constraints over a receding horizon [92]. The approach formulates the energy-management problem as a mixed-integer linear program that models storage dynamics, generator commitment, curtailment and load management. The methodology closes the loop on a physical experimental microgrid rather than remaining in simulation. The reported experimental results show that the approach is feasible on the physical system and that the schedules respect the

operating constraints under realistic forecast errors because the receding horizon repeatedly corrects the plan as new measurements arrive. The novelty is the experimentally validated connection between forecasting, optimization and physical actuation. Cominesi et al. extended this logic with a two-layer stochastic distributed MPC architecture, separating local decisions from network-level coordination under uncertainty [93]. In their approach, a slow layer performs scenario-based economic optimization while a faster shrinking-horizon layer refines the set points against short-term deviations. The methodology is a simulation case study of a microgrid with renewable generation, storage and controllable exchange evaluated against deterministic and single-layer alternatives. The reported results indicate close-to-ideal economic performance with satisfied constraints despite forecast uncertainty. The novelty is the explicit multi-rate separation between uncertainty-aware planning and fast correction. These approaches are particularly relevant because they already generate structured, constraint-aware plans. In a cognitive mini-grid, an agentic planner should therefore call an MPC or optimal-power-flow service rather than replace it. The agent supplies interpreted goals and contingency context, the optimizer supplies numerically feasible schedules and an independent assurance layer verifies the result.

The MAS literature supplies the organizational model. Dimeas and Hatziargyriou implemented a multi-agent microgrid controller in which local agents represent controllable entities and coordinate operation without requiring every decision to pass through one central process [11]. Their approach organizes production, consumption and coordination agents in a light hierarchy and lets them negotiate through a symmetrical-assignment market mechanism; the methodology is a prototype agent implementation exercised on a case-study microgrid with market-based dispatch. The reported results demonstrate that auction-based agent negotiation can settle local dispatch without a central decision process. The novelty is the first complete agent-organizational realization of microgrid control that includes its economic dimension. McArthur et al. separated MAS research into conceptual foundations and implementation technologies, emphasizing autonomy, social ability, messaging, ontologies and engineering tools [9,10]. The two-part methodology pairs a critical concept review with concrete guidance on standards, platforms and design patterns. The contribution is an engineering discipline for power-system MAS rather than an experimental result. Pipattanasomporn et al. showed how distributed smart-grid agents can be designed and implemented for local control tasks through an architecture in which control, distributed-energy-resource, user and database agents cooperate [103]. Their simulated demonstration isolates the customer system after an upstream outage and secures critical loads from local resources, so the reported outcome is a reproducible design template for distribution-level agent deployment. These papers established modularity and fault containment, but most communication remained syntactic: messages had agreed formats, yet the epistemic consequences of delivery, delay or conflicting interpretations have rarely been modeled.

Nunna and Doolla developed multi-agent distributed-energy-resource management for intelligent microgrids, associating agents with local objectives and coordination responsibilities and coupling energy management with incentive-based demand response across interconnected microgrids [12]. The methodology is a simulation-based agent implementation. The reported results indicate improved local balancing and reduced peak stress when consumer agents respond to incentives. The novelty is the integration of demand response into agent-level microgrid coordination. Morstyn et al. examined cooperative control of heterogeneous storage in a DC microgrid and demonstrated why distributed devices require coordinated state and power-sharing objectives rather than identical local rules [104]. Their methodology combines consensus-based information exchange with local feedback design and simulation studies; the reported results show state-of-charge

balancing and coordinated power sharing among non-identical units without a central controller, which makes the scheme a device-level template for asset agents. Zheng et al. used asynchronous consensus ADMM for multi-agent energy management, directly addressing the impracticality of perfectly synchronized updates [105]. The methodology decomposes the microgrid economic problem across agents that exchange variables under delays and asynchronous activation and the simulations report convergence to near-optimal schedules close to the synchronous benchmark; the novelty is the demonstration that distributed optimization can tolerate realistic communication imperfections. Howell et al. added holonic and semantic organization, showing how agents can form nested semi-autonomous groups and use shared information models for distributed energy resources [68]. Their methodology is an architectural review connecting semantic technologies with holonic control. The contribution is an organizational roadmap rather than a validated controller. Collectively, these works motivate the distributed-intelligence layer of the proposed architecture while also exposing open questions about semantic alignment, changing authority and verifiable commitments.

Table 14 reveals three successive shifts in the foundational literature. The first shift is from the microgrid as a controllable electrical entity to a rigorously separated hierarchy of primary, secondary and tertiary functions. Lasseter's formulation and the hierarchical frameworks of Guerrero et al. and Bidram and Davoudi establish the architectural principle that fast local regulation, restoration and slower economic coordination should not be collapsed into one controller [1,4,5]. This principle remains binding for cognitive minigrids, as argued in Section 3.1: an agentic planner can modify objectives or schedules, but it should not replace inverter protection or primary frequency response. The second shift is from fixed hierarchical control toward optimization-based coordination. MPC and distributed MPC make forecasts, constraints and future trajectories explicit, which provides a natural bridge to digital-twin validation. Their main limitation for the proposed architecture is not a lack of rigor but dependence on model fidelity, forecast quality and defined infeasibility procedures [92,93].

The third shift is organizational. Asset-oriented MAS [11,12], the MAS engineering discipline for power systems [9,10], cooperative control of distributed storage [104], consensus ADMM [105] and holonic architectures [68] distribute intelligence and remove a single supervisory bottleneck. However, the comparison shows that distribution is not equivalent to cognition. Most methods exchange numeric set points, dual variables or locally estimated states whose meaning is fixed by design. They rarely report ontology versions, source authority, confidence, acknowledgment status or the conditions under which a public signal becomes operational common knowledge. Furthermore, validation is concentrated in analytical models and simulation, with fewer studies examining adversarial communication, heterogeneous vendor semantics or the behavior of a distributed controller after a failed optimization or stale update.

The collective lesson is that the classical papers supply the non-negotiable physical foundation and the distributed papers supply scalable organization, but a cognitive extension must add explicit message semantics and evidence handling without disturbing their stability assumptions. A credible implementation should therefore reuse the mature controllers as certified tools or execution services. Semantic and agentic layers may propose changed references, coalition structures or schedules, while the established controllers retain responsibility for feasibility, transient response and safe fallback. This interpretation also explains why the proposed architecture is layered: it is designed to preserve the strongest result of each foundational family instead of replacing all of them with a single learned policy.

**Table 14.** Specific examination of foundational control, optimization and multi-agent papers.

| Paper | Problem and Control Level | Method/Architecture | Information Assumption | Validation Type | Main Contribution | Limitation and Cognitive-Grid Relevance |
|---|---|---|---|---|---|---|
| Lasseter [1] | Definition and operation of a controllable microgrid entity | Local source control, islanding and system-level coordination concept | Local measurements and supervisory coordination | Conceptual/system studies | Established the bounded microgrid as a controllable electrical entity | Does not model semantic knowledge or AI; supplies the physical contract that cognitive layers must respect |
| Guerrero et al. [4] | Power sharing, restoration and tertiary management | Three-level hierarchical control for AC/DC systems | Explicit control hierarchy and time-scale separation | Analytical and representative case studies | Generalized primary, secondary and tertiary architecture | Fixed roles and messages; becomes the authority/time-scale backbone of cognitive control |
| Bidram and Davoudi [5] | Classification of microgrid controller functions | Centralized and distributed hierarchy | Communication needs vary by layer | Survey/analytical synthesis | Clarified controller responsibilities and structures | Does not capture agent beliefs; useful for defining prohibited cross-layer actions |
| Parisio et al. [92] | Constrained economic operation | Receding-horizon MPC | Forecasts and explicit models are available | Experimental microgrid optimization | Demonstrated constraint-aware predictive scheduling | Model and forecast dependence; suitable as a verified tool called by an agent |
| Cominesi et al. [93] | Distributed operation under uncertainty | Two-layer stochastic distributed MPC | Coordinating variables exchanged among local optimizers | Simulation-based evaluation | Connects uncertainty management with distributed optimization | Communication semantics and cyber failures are not central; forms a numerical coordination service |
| Dimeas and Hatziargyriou [11] | Decentralized microgrid operation | Asset-oriented MAS | Agents exchange operational messages | Prototype/simulation | Early demonstration of autonomous agent organization | Semantic equivalence and common knowledge are implicit rather than verified |
| McArthur et al. [9,10] | MAS concepts and engineering practice for power systems | Agent models, communication, standards and tools | Shared protocols and ontologies | Two-part survey | Established a rigorous MAS vocabulary for power engineering | Predates current MARL and LLM agents; remains essential for role and protocol design |
| Nunna and Doolla [12] | DER management in intelligent microgrids | Distributed agents with local and global objectives | Neighbor or supervisory information | Simulation | Demonstrated modular DER coordination | Limited treatment of uncertain semantics and adversarial messages |
| Morstyn et al. [104] | Heterogeneous storage coordination | Cooperative multi-agent control | Local and neighboring storage states | Analytical/simulation studies | Coordinates non-identical storage assets | Assumes a trusted control channel; can become a device–agent baseline |
| Zheng et al. [105] | Asynchronous energy management | Consensus ADMM in an MAS | Delayed or asynchronous variable exchange | Simulation | Removes strict synchronization requirement | Optimizes numeric variables but not shared meaning or authority |
| Howell et al. [68] | Organization of highly distributed energy resources | Semantic and holonic MAS | Shared semantic models among nested agent groups | Architectural review/proposal | Links semantics with scalable organizational structure | Requires operational ontologies and safety interfaces before closed-loop deployment |

### *4.2. Reinforcement Learning, MARL and Common-Knowledge Coordination*

Learning-based methods form the first family that attempts to relax the model dependence of the foundations examined above. The study of Kuznetsova et al. was among the early works that formulated microgrid energy management as an RL problem in which sequential decisions are learned from interactions rather than solved exclusively from a known model [161]. The approach applies tabular Q-learning to battery scheduling in a microgrid with wind generation, using a two-step-ahead decision formulation. The methodology is a scenario-based simulation study of repeated daily operation. The reported results show improved battery utilization and better attainment of the customer's reliability goals across the simulated scenarios. The novelty is the demonstration that a learned sequential policy can encode operational experience under generation uncertainty. The importance of the work lies in its treatment of uncertainty and repeated operation. Its limitation is shared by many early RL studies: the environment and reward define what the controller regards as correct, while electrical safety is usually enforced indirectly through penalties or through a simplified simulator.

Foruzan et al. used RL for distributed energy management and showed how local decision makers can learn coordinated economic behavior [119]. The approach assigns learning agents to generation units, storage and customers inside a microgrid energy market. The methodology is a multi-agent Q-learning simulation over repeated market interactions. The reported results show convergence toward near-optimal bidding and dispatch strategies, together with scalability across microgrid sizes; the novelty is a fully distributed learned management scheme in which no agent requires the internal models of the others. Zeng et al. combined approximate dynamic programming with deep recurrent learning to represent temporal dependencies in microgrid operation [120]. Their methodology approximates the value function with a recurrent network trained over stochastic operating scenarios. The reported results approach the offline optimization benchmark while remaining executable online; the novelty is the explicit modeling of memory in the learned energy-management decision process. Ji et al. studied real-time energy management using deep RL, demonstrating the attraction of direct state-to-action policies for online operation [121]. The methodology trains a deep Q-network on stochastic renewable, load and price scenarios. The reported results show lower expected operating costs than myopic rules while approaching the cost of an idealized deterministic schedule, without requiring explicit forecasts at decision time. Du and Li extended model-free learning to a multi-microgrid setting, where coordination must account for exchanges among several local systems [122]. In their approach, a distribution-level operator learns retail decisions with a model-free method while a deep network surrogate represents the aggregated response of each microgrid, so internal models remain private; the reported simulation results indicate reduced system-level peaks and costs. The novelty is coordination across electrical boundaries without internal-model disclosure. These papers progressively enlarge the state and decision spaces, but they also increase the difficulty of tracing why a particular action was selected.

Goh et al. examined multistage reward-function design and showed that reported RL performance can be strongly influenced by how objectives and constraint violations are encoded [162]. Their methodology evaluates alternative reward stages on the same microgrid energy-management task. The reported results show that a multistage design improves convergence behavior and constraint adherence relative to single-stage rewards; the principal contribution is methodological. Namely, that conclusions about "RL performance" are partly conclusions about reward engineering. Ye et al. explicitly addressed safe deep reinforcement learning (DRL) for microgrid energy management by integrating learning with operational limits and a spatiotemporal perception of the distribution

network state [123]. The methodology couples a safety-constrained learning formulation with a spatial–temporal state representation of the distribution network and evaluates it in closed-loop distribution-network simulations; operating costs competitive with unconstrained DRL, together with substantially reduced constraint violations, are reported. The novelty is the treatment of safety as an explicit optimization structure rather than a penalty weight. Lee et al. proposed a newer DRL-based EMS architecture that continues the trend toward integrated learned decision pipelines, reporting improved economic and operational indicators against the compared baselines in simulation [163]. The key lesson is that "safe RL" should not be inferred from a low observed violation rate. A deployment claim requires a defined safe set, a mechanism that prevents or corrects unsafe actions and tests under distribution shifts and rare contingencies. Learned policies should be layered with deterministic fallback control: the learning policy proposes economically or operationally useful actions, while a shield, verifier or backup controller preserves invariants.

Beyond the power domain, the general MARL literature supplies the coordination mechanisms on which these studies build. Multi-agent deep deterministic policy gradient (MADDPG) trains centralized critics with decentralized actors [131]. Value-decomposition networks (VDNs) and QMIX factorize a joint value function for decentralized execution [133,134]. Counterfactual multi-agent policy gradients (COMA) address credit assignment with counterfactual baselines [132]. MACKRL conditions group policies on common knowledge and local policies on individual information [15]. Safe-learning research adds shielding [150], Lyapunov-constrained policy updates [164] and the broader safe-RL taxonomy [64]. These works are cited for the ideas they contribute; because their benchmarks are not electrical, they are not examined in Table 15.

In a mini-grid, common knowledge could arise from synchronized frequency measurements, authenticated broadcasts, acknowledged commitments or a replicated event log. However, observing the same physical variable is not always equivalent to having common knowledge: sensor errors, local filtering and communication failures may cause different interpretations. The cognitive mini-grid therefore requires explicit acknowledgment and validity semantics when a coordinated action depends on a shared fact.

The learning studies in Table 15 differ along two decisive dimensions: where coordination occurs and how safety is represented. Early and centralized RL studies learn an energy-management policy from repeated interaction, usually with feasibility embedded in the environment, penalties or the action space [121,161]. Distributed RL [119] and multi-microgrid studies [122] move decisions toward local controllers and inter-grid exchanges, improving scalability and adaptation but also increasing non-stationarity because each learning agent changes the environment observed by the others. The recurrent approximate-dynamic-programming approach instead retains a single EMS-level decision process while enlarging its temporal memory [120]. Reward-design studies demonstrate that apparently superior policies may be artifacts of the selected reward decomposition, which means that cost or cumulative reward alone cannot establish engineering superiority [162]. The safe-DRL approach improves this position by making admissibility or risk part of the learning architecture rather than a post hoc observation [123].

Coordination under partial observation exposes a separate issue. For mini-grids, the distinction between group policies conditioned on common knowledge and local policies conditioned on individual information is operationally important. A group action such as coordinated islanding, black start or reserve commitment should not be triggered merely because several agents possess similar local observations; the relevant state and authorization must be sufficiently shared and acknowledged. The general MARL literature demonstrates the algorithmic value of shared information but does not, by itself, specify

how electrical agents establish that information under packet loss, inconsistent topology models or compromised broadcasts.

**Table 15.** Specific examination of learning and coordination studies relevant to cognitive mini-grids.

| Paper/Method | Decision Problem | Learning Mechanism | Coordination Model | Safety Treatment | Evidence Contributed | Main Unresolved Issue |
|---|---|---|---|---|---|---|
| Kuznetsova et al. [161] | Sequential energy management under uncertainty | RL policy learned from repeated operation | Primarily local/EMS decision process | Constraints represented through environment/reward | Established RL as a viable microgrid EMS formulation | Limited safety evidence and transfer across operating conditions |
| Foruzan et al. [119] | Distributed economic management | Distributed RL | Local decision makers coordinate energy use | Operational constraints in problem design | Demonstrated distributed learned resource management | Convergence and stability under communication failures require deeper analysis |
| Zeng et al. [120] | Dynamic energy management with temporal dependencies | Approximate dynamic programming and recurrent modeling | EMS-level state/action process | Feasibility embedded in the model | Captures sequential uncertainty and memory | Limited interpretability and out-of-distribution guarantees |
| Ji et al. [121] | Online microgrid scheduling | Deep RL | Central EMS policy | Penalty/constraint formulation | Shows direct online policy inference | Does not, by itself, provide certifiable safety |
| Du and Li [122] | Coordination among multiple microgrids | Deep neural network plus model-free RL | Inter-microgrid exchanges | System constraints in simulation | Extends learning beyond one electrical boundary | Shared-state, privacy and strategic behavior assumptions need scrutiny |
| Goh et al. [162] | Effect of reward design on EMS behavior | Multistage reward analysis | Depends on selected RL architecture | Constraints encoded through reward stages | Demonstrates sensitivity of conclusions to reward engineering | Reward compliance is not equivalent to invariant enforcement |
| Ye et al. [123] | Energy management with operational safety | Safe DRL | EMS-level controller | Explicit safety mechanism/limits | Moves evaluation from reward alone toward safe operation | Rare-event, model-mismatch and real-time assurance remain open |

Accordingly, Table 15 supports a bounded role for learned policies. RL is well suited to scheduling, adaptation, forecasting-assisted decisions and candidate action selection under uncertainty. It is less suitable as the sole source of fast protection or as the final authority for high-consequence switching unless a separate safety mechanism proves the admissibility of the action. Future comparisons should therefore report not only reward and mean operating cost but rare-event violations, distribution shift, asynchronous observations, unsafe proposals rejected by a shield, verifier latency and policy-update governance. This evidence would distinguish a high-scoring simulator policy from a learning component that can be responsibly integrated into the cognitive architecture.

### *4.3. Semantic Communication, Task-Oriented Messaging and Semantic Control*

Every learning or coordination method examined so far assumes that agents can exchange the information their policies require; semantic communication addresses what should be exchanged and how its meaning is preserved. Outside the power domain, the semantic-communication literature establishes the idea itself: semantic information theory [19], the 6G view of semantic and goal-oriented communication [20,21], learned end-to-end semantic transmission [23], task-oriented multi-user semantics [165] and significance- and age-aware messaging [22,166]. These works supply the vocabulary of semantic con-

trol adopted in Section 3.2, but they were not evaluated on electrical tasks, so only the power-specific study is examined in Table 16.

Diao et al. examined neuromorphic event-driven semantic communication in the microgrid domain [27]; the study is summarized in Table 16. The approach embeds a spiking neural network (SNN) at each converter node. The SNNs are trained collaboratively online using only the power exchanges between the nodes so that remote information is inferred from local electrical measurements without a dedicated communication channel, while an event-driven selective process collects the sparse data used for training. The methodology validates the scheme in simulation across different microgrid topologies and components and in a laboratory test with two DC/DC converters subjected to load-change and line-outage events. The reported results indicate reliable coordination performance without the exogenous data paths of a conventional communication layer. The novelty is the direct co-design of semantic representation, event-driven signaling and microgrid control. The work is important for three reasons. It treats communication as part of the physical coordination problem, uses event-driven representations that are compatible with sparse grid disturbances and targets low-overhead embedded operation. Its evidence is a meaningful proof of concept, yet it does not establish a general semantic standard for heterogeneous devices or demonstrate that compressed messages preserve every fact needed for protection and formal verification.

Semantic control also changes the definition of common knowledge. A broadcast is not sufficient when coordinated action depends on every agent knowing that every other relevant agent received and understood the same instruction. A reliable semantic-control protocol therefore needs message identity, source authority, context, confidence, validity horizon, required acknowledgments and a commitment state. These fields can be added above IEC 61850, the Common Information Model (CIM), IEEE 2030.5 or MQTT without replacing existing device and service models.

Conventional semantic metrics are insufficient for a power mini-grid. Sentence similarity, classification accuracy or task success can remain high, even when the omitted information is precisely what a controller requires to verify safety. A message such as "storage available" may preserve the general meaning while concealing the maximum power, duration, reserve policy, state of health or validity interval. The study of Diao et al. is especially important because it moves the evaluation into a microgrid control context and associates sparse event-driven representations with electrical coordination [27]. Nevertheless, this emerging line still needs broader tests covering packet loss, delayed acknowledgments, ontology mismatch, malicious semantic substitutions and transient electrical consequences.

The principal comparative conclusion is that semantic compression must be safety-aware and reversible. The receiver should be able to request raw evidence when confidence is low, when the action consequence is high or when the semantic codebook version is inconsistent. The transmitted object should include not only a compact intent but provenance, confidence, physical scope, units and validity horizon. Evaluation should compare raw-data streaming, conventional event triggering and semantic messaging under the same communication budget and disturbance scenarios. In this way, the benefit claimed by semantic control can be separated into bandwidth reduction, faster decision formation and improved coordination, while any loss of observability or verification coverage is measured explicitly.

### *4.4. LLM-Based Agents, Tool Use and Supervisory Autonomy*

Semantic messages must be interpreted, related to operating rules and turned into proposals by a deliberative component; LLM-based agents are the most recent candidates for that role. The general agent literature provides the mechanisms that the grid-specific

systems reuse: RAG grounds generated text in retrieved documents [30], ReAct interleaves reasoning steps with tool actions [31], Toolformer learns when to call external tools [32] and AutoGen and MetaGPT organize role-specialized multi-agent workflows [44,45]. These mechanisms were validated on language and software benchmarks without electrical deadlines, so they are referenced for their design ideas and are not examined in Table 17.

Grid-Agent applies these mechanisms directly to power-grid control [33]. A planning agent produces coordinated action sequences and interacts with numerical power-flow tools, while a validation agent evaluates the proposal in a sandbox and supports rollback. The methodology exercises the framework on violation-mitigation scenarios in standard IEEE and CIGRE networks. Successful mitigation of injected violations, including multiple simultaneous violations, is reported across six LLM configurations of the same architecture, with performance varying by the selected model. These experiments demonstrate that semantic reasoning can be combined with numerical analysis. Formal optimality is not established by the study. The paper's central contribution is not that an LLM replaces optimization; it is the modular separation of planning and validation. The main limitation for mini-grid deployment is evidence maturity: simulation success does not establish bounded latency, cyber isolation, prompt robustness, protection compatibility or authority management on a live system.

X-GridAgent focuses on assisting in power-system analysis through a hierarchical planning, coordination and action architecture [34]. Natural-language interaction is combined with domain tools and structured retrieval, including mechanisms for prompt refinement and schema-adaptive RAG. The reported evaluation covers power-flow, optimal-power-flow and contingency-style queries. Previously unseen analysis tasks were decomposed and completed from conversational input in the evaluated cases. The layered translation from open-ended language to workflows executed through numerical tools is the architectural contribution. This makes it an operator-support architecture rather than a millisecond controller. Its strength is extensibility and the ability to decompose previously unseen analysis requests. Its risk is that tool selection, parameter construction and interpretation remain model-mediated; the system therefore needs typed interfaces, provenance records and independent checks of every numerical result.

The semantic risk-aware VPP direction uses language models to interpret unstructured warnings or contextual information and translate them into optimization constraints or pre-activation decisions [167]. The methodology introduces a semantic demand–response pre-activation constraint into a virtual power plant dispatch model, where interpreted early-warning text modifies reserves and flexible-load commitments before the risk materializes. Improved dispatch outcomes relative to the deterministic baseline are reported in a seven-day controlled stress test with synthetic alert messages; contextual language is incorporated through scenario generation, penalties and dispatch constraints. This is a promising example of semantic control because natural-language information becomes a formally constrained input to an optimization model. The critical boundary is the translation step: ambiguous text must not silently become a hard operational constraint. Confidence, source credibility, conflict resolution and human approval should be explicit when the interpreted information changes reserve, curtailment or market commitments.

**Table 16.** Specific examination of semantic and task-oriented communication research.

| Paper | Semantic Objective | Representation and Optimization | Receiver Task | Validation Context | Contribution to This Survey | Gap for Mini-Grid Deployment |
|---|---|---|---|---|---|---|
| Diao et al. [27] | Replace exogenous data exchange with inference from power flows | Spiking/neuromorphic event representation | Distributed microgrid coordination | Simulation across microgrid topologies and a two-converter laboratory test | Directly links semantic communication and mini-grid control | Generalization, interoperability, cyber resilience and hardware trials are limited |

**Table 17.** Specific examination of LLM and agentic systems relevant to power and mini-grid operation.

| System/Paper | Agent Roles | Grounding and Tools | Output Authority | Evidence | Important Contribution | Required Extension for Cognitive Mini-Grids |
|---|---|---|---|---|---|---|
| Grid-Agent [33] | Planner and validation agents | Power-flow solvers and sandbox evaluation | Candidate corrective control with validation/rollback | IEEE and CIGRE simulation cases | Demonstrates composition of semantic planning and numerical precision | Requires real-time tests, cyber isolation, protection studies and formal authority limits |
| X-GridAgent [34] | Planning, coordination and action layers | Domain tools, structured databases and hybrid RAG | Analysis assistance | Multiple grid-analysis queries and cases | Automates interpretable workflows through natural language | Requires reproducible model versions, typed tool calls and safety cases for operational use |
| Semantic virtual power plant (VPP) framework [167] | Semantic interpreter plus optimizer | Textual warnings/context linked to dispatch model | Changes optimization inputs or pre-activation | Emerging optimization studies | Connects unstructured context with constrained energy decisions | Source credibility, ambiguity, approval and adversarial text must be addressed |
| Agentic smart-grid synthesis [35] | Reviewed architectures span multi-agent, RL, digital-twin and physics-based control | Not a system; proposes evaluation criteria of stability, safety, interpretability, certification readiness and interoperability | Ranges from advice to autonomous control across the reviewed systems | Survey of reviewed prototypes | Organizes a rapidly emerging field and its safety concerns | Mini-grid-specific latency, offline operation and proof-carrying execution need deeper study |

The wider agentic-AI review literature identifies planning, memory, tool use and autonomous execution as defining capabilities. Wang et al. structured LLM-based autonomous agents around profile, memory, planning and action modules [168] and Xi et al. surveyed the rise and potential of LLM-based agents, including their communication and societal risks [46]. The smart-grid-specific review by Kiasari and Aly surveys agentic architectures, multi-agent and RL control, digital-twin optimization and physics-based control [35]. The review proposes an evaluation framework based on stability, safety, interpretability, certification readiness and interoperability with grid codes; its principal contribution is a consolidated synthesis and evaluation framework for certifiable agentic grid control rather than a new system [35]. For cognitive mini-grids, these capabilities must be assigned to the bounded roles specified in Section 3.3. A planner may generate candidate actions, an analyst may summarize incidents, a market agent may negotiate within approved limits and a validator may reject unsafe proposals. No language agent should have unrestricted write access to protection settings, inverter firmware or switching equipment. The architecture must also remain functional when cloud models or external networks are unavailable.

The entries in Table 17 provide direct domain relevance. Grid-Agent combines a planning agent with numerical validation and sandbox evaluation, while X-GridAgent organizes planning, coordination and action layers around grid-analysis tools and structured retrieval [33,34]. The semantic VPP framework further shows how textual warnings or contextual information can alter the inputs or pre-activation constraints of an optimization problem [167]. Compared with generic chat-based demonstrations, these systems make an important methodological advance: the LLM does not calculate the final power-flow result from language alone but delegates precise computation to domain tools. Even so, the validation is mainly simulation-based and the remaining questions concern model versioning, real-time performance, stale retrieval, tool failure, adversarial prompts, cross-agent impersonation and the formal limit of action authority.

The comparison supports a clear deployment rule. LLM or SLM agents should initially be treated as planners, analysts, coordinators and interfaces. Their outputs should be typed objects that contain goals, selected tools, parameters, assumptions and evidence links. A separate deterministic service should validate units, topology, state freshness, power-flow feasibility, stability constraints, access rights and rollback readiness. Only the certified result should reach an actuator. Evaluation should report tool-call correctness, unsupported statements, unsafe proposals, the validation catch rate, total latency and recovery after missing or contradictory data. Under this design, agentic AI adds flexibility and explanation without becoming an unbounded control path.

### *4.5. Digital Twins, Neuro-Symbolic Assurance and Federated Foundation Models*

Keeping that control path bounded requires a place where proposals can be evaluated and a mechanism by which they are admitted or rejected. Digital-twin research provides the environment in which a proposed action can be tested before it reaches the physical mini-grid, which is the role assigned to twins in Section 3.4. The general twin literature defines the concept [61,140]; distinguishes an offline digital model, a data-connected digital shadow and a bidirectionally coupled digital twin [139]; and emphasizes synchronization, fidelity, lifecycle integration and uncertainty as open challenges [138,141]. These domain-independent definitions are adopted here without separate examination. Recent power-system digital-twin studies move the concept toward operational use [60,142]. In the deployment-oriented account, lessons are distilled from two laboratory-scale live twins of DC microgrids that share controllers and real-time data streams with their physical counterparts; fidelity standards, multi-physics modeling and real-time synchronization are

reported as the principal challenges [142]. In the smart-grid introduction, a twin is defined as a real-time, high-fidelity representation that controls its physical counterpart through computational reflection [60]. In this survey, latency guarantees and model fidelity are therefore treated as explicit requirements rather than optional refinements. For a cognitive mini-grid, topology, parameters, asset states, communication state and uncertainty should be maintained by the twin. Counterfactual questions should be answered about whether a proposed switching sequence causes overload, whether a delayed battery response compromises frequency recovery and which assumption caused a plan to fail. A causal twin is therefore more valuable than a purely predictive surrogate when the system encounters a disturbance outside the training distribution.

Neuro-symbolic AI addresses a related assurance need. Uncertain states can be inferred, anomalies can be detected and actions can be proposed by neural components, while grid rules, switching interlocks, authorization policies and explanation traces can be encoded by symbolic components. The foundations of neural-symbolic learning systems [169] and the case for a third wave of neuro-symbolic AI [47] supply the general framework, while a grid-oriented preprint applies it to autonomous grid operation [62]. In that preprint, which, under the source policy stated in Section 2.1, is treated as evidence of an emerging direction only, control actions are proposed by a trainable actor–critic policy, while rules are enforced and traceable justifications are produced by a logic-based module. Better performance than purely neural and purely rule-based controllers is reported in simulations on a distribution-grid testbed, together with traceable decision records [62]. The principal contribution is the treatment of explainable autonomous operation as a hybrid reasoning problem rather than as post hoc visualization. Evidence maturity remains a limitation because symbolic rules may be incomplete and neural outputs may be miscalibrated. Conflict handling and deterministic fallback are therefore required by the architecture.

Proof-carrying code offers a useful assurance analogy: untrusted code is accepted only when accompanied by evidence that a specified policy holds [73]. In the original formulation, a machine-checkable proof is attached to the code and acceptance is decided by a small trusted verifier rather than by trust in the producer [73]. A proof-carrying agent action similarly includes the proposed command, initial-state assumptions, model and tool versions, predicted consequences, constraint margins, authorization and a compact certificate. The verifier is intentionally simpler than the planner. Model error is not eliminated, but the execution decision is shifted from trust in the agent's narrative to verification of explicit claims.

Federated learning addresses the fact that mini-grid operators may benefit from shared models without transferring raw customer or operational data. The canonical algorithms are FedAvg, which averages locally trained models under non-identically distributed data [16], and FedProx, which adds a proximal term for heterogeneous clients [170]; the FL literature further documents privacy, communication, poisoning and participation challenges [17,124]. A mini-grid federation adds topology differences, seasonal regimes, hardware diversity and regulatory boundaries. Consequently, a globally averaged model may be inferior to a personalized or clustered model.

Power-grid foundation-model research is beginning to explore reusable representations and general-purpose operational assistance [52]. Federated foundation models extend this idea across organizations [125], while trustworthiness layers seek calibrated uncertainty or conformal evidence for high-consequence decisions [54]. These are future-oriented contributions rather than deployment-ready systems. Their value to the survey is that they reveal a possible shared cognitive substrate; their limitation is the unresolved cost of training, update governance, data provenance, catastrophic forgetting, security and physical certification.

The studies in Table 18 occupy different parts of an assurance pipeline and should not be treated as interchangeable "digital twin" solutions. The distinction between an offline digital model, a data-connected shadow and a bidirectionally coupled twin is essential for mini-grids because a planning agent may produce a valid result in an offline model while the physical topology, storage state or protection status has already changed. Power-grid twin studies bring the concept closer to operational monitoring and control, but the evidence remains heterogeneous and standard interfaces, causal validation and field demonstrations are still limited [60,142].

The neuro-symbolic entry addresses trust from a different direction. Neuro-symbolic grid AI combines learned adaptation with explicit rules, improving traceability and allowing known constraints to override a neural proposal [62]. Proof-carrying code contributes the stronger architectural idea that an untrusted producer can submit a proposal together with machine-checkable evidence, after which a small trusted verifier decides whether to accept it [73]. In a cognitive mini-grid, the digital twin can generate predicted consequences, the symbolic layer can encode policy and physical invariants and the proof-carrying interface can bind the proposal to its assumptions and validity horizon. The principal open issue is the construction of certificate languages and verifiers that are expressive enough for grid constraints but fast and simple enough to be independently trusted.

Foundation-model studies add cross-site intelligence rather than direct operational assurance. Emerging grid and federated foundation models aim to reuse multimodal representations across sites and tasks [52,125]. Their benefits are strongest for forecasting, diagnosis, retrieval and representation learning. Their risks include poisoning, drift across non-identically distributed sites, excessive communication, uncertain model provenance and unsafe generalization to closed-loop decisions. Table 18 therefore suggests a division of labor: federated models provide hypotheses and forecasts, twins evaluate consequences, symbolic policies express constraints and a proof verifier controls execution. No single component should certify its own output.

**Table 18.** Specific examination of twins, assurance and distributed foundation-model research.

| Paper/Stream | Primary Purpose | Model or Assurance Mechanism | Connection to the Physical System | Evidence Maturity | Contribution to Cognitive Mini-Grids | Principal Limitation |
|---|---|---|---|---|---|---|
| Power-grid twins [60,142] | Apply twins to monitoring, analysis and grid operation | Power-system models linked with operational data | Intended closed-loop or operator-support connection | Laboratory-scale live twins and an introductory book chapter | Brings twin concepts into grid workflows | Standard interfaces, causal validation and field evidence are limited |
| Neuro-symbolic grid AI [62] | Combine adaptation with explicit rules and explanations | Neural proposals plus symbolic constraints/reasoning | Intended supervisory decision support/control | Emerging preprint/ prototype | Provides interpretable safety logic around learned behavior | Rule completeness, conflict handling and independent replication are needed |
| Grid/federated foundation models [52,125] | Reuse representations across tasks and sites | Pretraining, adaptation and federated collaboration | Primarily analytical/model layer | Early research | Could supply shared multimodal intelligence across mini-grids | Compute cost, data governance and certification are unresolved |
| Trust layer for grid foundation models [54] | Add calibrated trustworthiness to foundation-model outputs | Uncertainty/ conformal-style assurance layer | Decision-support interface | Emerging preprint | Moves beyond raw prediction accuracy | Coverage under distribution shift and closed-loop consequences need validation |

### *4.6. DAI/BDIx Agent Frameworks and Flexible Power-Interface Control for Power Grids*

Two complementary research lines connect the preceding clusters to deployable mini-grid intelligence and are therefore examined with the same template. The first is the DAI framework with BDIx agents developed by Ioannou et al. [149]. The second is the flexible power-interface and power-flow assignment line developed by Javaid et al. [147,148,171]. The first line informs the distributed decision and learning layers of the proposed architecture. The second provides power-flow control methods and balanceability conditions within explicitly modeled device and connection limits. Their integration into the complete cognitive architecture is proposed here. Table 19 compares the individual approaches and their associated papers.

The DAI/BDIx line originated in wireless networking, where a BDIx agent embedded in every device selects its device-to-device (D2D) transmission mode autonomously [172]. The framework was later generalized through a DAI framework with machine learning for D2D communication in fifth-generation (5G) and 6G networks [173], applied to dynamic D2D communication [174] and consolidated in a monograph on distributed AI for 5G/6G communications [175]. These studies are cited for the agent architecture only. Because their evaluation is not electrical, they are not examined in Table 19. The framework was transferred to power systems in a nano-grid setting [149]. In this approach, BDIx agents are attached to fluctuating generators, fluctuating loads and storage devices; maintain local beliefs about the power state; and negotiate balancing actions so that supply–demand equilibrium is preserved without a central controller. The methodology is a simulation campaign over fluctuating photovoltaic generation and stochastic consumption in which the distributed agent organization is compared with centralized management across operating scenarios. Sustained power balance under the examined fluctuations is reported together with faster local reactions and lower coordination overhead than the centralized alternative; the reported values, the baseline, the system size and the platform of this and the other studies of the two lines are collected in Table 20, together with the corresponding entries for the semantic-communication study of Section 4.3 and the predictive-control experiment of Section 4.1, so that the evidence class assigned to each study in Table 21 can be read against the outcome it rests on. The novelty is the demonstration that the DAI framework can be instantiated as an energy-management organization in which every electrical asset hosts an autonomous agent. A companion conference study embedded forecasting into the same agent organization so that predicted power trajectories support mini-grid stability decisions [176]. The relevance to this survey is architectural: the BDIx agent supplies exactly the belief maintenance, desire arbitration and intention execution cycle that Section 3 assigns to the distributed-intelligence layer.

The learning-oriented continuation addresses state-of-charge balancing across heterogeneous battery systems [177]. The hierarchical predictive-adaptive control framework combines a long-horizon predictive engine with a fast adaptive layer. The predictive engine is implemented as a federated Transformer that produces multi-horizon probabilistic net-load forecasts across collaborating mini-grids, whereas in the adaptive layer, a DRL policy corrects the balancing decisions online. The framework is evaluated with thirteen alternative controllers under heterogeneous battery conditions and stochastic generation, with sensitivity analysis of the reward weights and ablation of forecasting and reward-shaping components. Improvements in state-of-charge balance and operating behavior are reported for the evaluated scenarios: in a MATLAB R2025a digital-twin environment and in a comparison of fourteen controllers, including the proposed controller, spanning rule-based, model-predictive and deep-reinforcement-learning families, the framework is reported to reach the lowest state-of-charge variance (described as essentially zero) and the lowest bus-voltage variance of the fourteen at near-minimal operating cost and

moderate energy throughput over a 24 h base case and three-day runs, with robustness shown under high net-load volatility and under communication impairments of the forecast channel. The architectural distinction relevant to this survey is the separation of federated forecasting from adaptive control within one learning stack. The forecasting branch was extended through a physics-guided self-supervised graph temporal Transformer that predicts electricity inconsistencies in mini-grids by exploiting the electrical graph structure and physics-consistent pretraining objectives, with reported gains over purely data-driven forecasting baselines [178]: on a six-node proxy graph derived from the University of California, Irvine (UCI) household consumption dataset with synthetically injected inconsistency episodes, it reaches an area under the receiver operating characteristic curve (AUC-ROC) of 0.8959, an F1 score of 0.8307 and a false-alarm rate of 0.41 percent, an F1-score 2.2 percent higher than the strongest recurrent baseline and up to 15.2 percent higher than the long short-term memory (LSTM) baseline, with a false-alarm rate about 52 percent lower than the LSTM baseline. On the IEEE 13-node test feeder simulated in OpenDSS, it reaches an area under the curve of 0.9016 and an F1 score of 0.8361 against eight baselines.

The flexible power-interface line approaches the same physical problem from the opposite direction: instead of learning behavior, it makes power flows individually designable and formally checkable. The power-flow coloring concept assigns a unique identifier to each power flow between a specific source and a specific load so that multiple simultaneous flow patterns can be designed with respect to availability, cost and emissions [171]. The implementation study realized the concept over a nano-grid with fluctuating loads through a cooperative distributed control method in which a master–slave role assignment among sources and a time-slot-based feedback loop absorb the fluctuations [147]. The methodology includes a physical implementation. Maintenance of the designated flow patterns and stable voltage under load fluctuations is reported; the extension to simultaneously fluctuating sources and loads compiles a power-flow-pattern matrix into per-device power ratios and demonstrates the resulting operation [171]. The novelty of the line is the transformation of aggregate power balance into individually accountable source-to-load flows, which anticipates the per-commitment accountability that agentic coordination requires.

The balanceability branch of the same line replaces empirical tuning with formal feasibility. A system characterization that states safe operation conditions among power generators, loads and storage devices was derived so that the balanceability of a configuration with given fluctuation ranges can be decided [148]. An efficient testing scheme reduces the computational effort of this decision compared with exhaustive checking [179] and the storage-sizing study derives the system condition under which power balancing is guaranteed and computes optimized storage sizes that satisfy it [180]. The methodology across these papers is analytical characterization supported by numerical demonstrations. The reported results are checkable conditions and minimum storage capacities rather than empirical performance curves; the novelty is precisely this checkability, which makes the conditions natural candidate predicates for the proof-carrying actions developed in Section 7.1. The joint continuation applies multiple-load power-flow assignment to loss mitigation under incremental load variations while user comfort is preserved [181]. Four logical interconnection scenarios among generators, loads and storage are evaluated with three control algorithms based on total demand, adaptive demand and grid support, using measured photovoltaic and fuel-cell data across seasons. Reduced storage-related losses are reported as interconnection richness increases, full daily demand satisfaction is achieved when grid support is available and critical loads are prioritized through the adaptive policy when grid support is limited.

**Table 19.** Specific examination of the DAI/BDIx agent approaches of Ioannou et al. and the flexible power-interface approaches of Javaid et al. for power grids.

| Paper | Problem and Setting | Approach and Method | Validation | Results in Brief | Novelty and Contribution | Role in the Cognitive Mini-Grid |
|---|---|---|---|---|---|---|
| Ioannou et al. [149] | Power balance of a nano-grid with fluctuating sources, loads and storage | BDIx agents per electrical asset negotiating balancing actions | Simulation campaign against centralized management across scenarios | Sustained supply–demand balance with faster local reactions and lower coordination overhead | Instantiation of DAI for distributed power management and control | Direct instantiation of the distributed-intelligence layer on grid assets |
| Ioannou et al. [176] | Forecast support for mini-grid stability | Agent-embedded power forecasting inside the DAI organization | Conference-scale simulation demonstration | Predicted trajectories improved agent balancing decisions | Coupling of forecasting with agent deliberation | Perception-to-belief pathway of the agent cycle |
| Ioannou et al. [177] | State-of-charge balancing across heterogeneous battery systems in mini-grids | Hierarchical control: federated Transformer forecasts plus a DRL adaptive layer | Comparison with thirteen alternative controllers; sensitivity and ablation studies | Improved balance and operation in the evaluated scenarios | Separates federated forecasting from adaptive control | Combined learning and coordination layers with time-scale separation |
| Ioannou et al. [178] | Forecasting electricity inconsistencies in mini-grids | Physics-guided self-supervised graph temporal Transformer | Comparative evaluation against data-driven forecasting baselines | Improved inconsistency forecasting through physics-consistent pretraining | Graph-and-physics-aware self-supervision for grid time series | Twin-adjacent predictive service for agents and validators |
| Javaid et al. [147,171] | Individually designable power flows over nano-grids with fluctuating loads and sources | Power-flow coloring with unique per-flow identifiers; master–slave roles, time-slot feedback and pattern compilation | Physical implementation and demonstrations | Maintained designated flow patterns and stable voltage under fluctuations | Aggregate balance becomes accountable source-to-load flows | Verified execution substrate for per-commitment accountability |
| Javaid et al. [148,179,180] | Safe operation, balanceability and storage sizing with fluctuating and controllable devices | Formal balanceability characterization, efficient testing scheme and storage-size optimization | Analytical conditions with numerical demonstrations | Checkable safe-operation conditions and minimum storage sizes that guarantee balance | Formal, checkable feasibility replaces empirical tuning | Certificate predicates and planning-time assurance input for the validator |
| Muhammad et al. [181] | Loss mitigation under incremental load variations in distributed power-flow systems | Multiple-load power-flow assignment with total-demand, adaptive-demand and grid-based control | Four interconnection scenarios with measured seasonal photovoltaic and fuel-cell data | Losses reduced with richer interconnection; daily demand met with grid support; critical loads prioritized otherwise | Loss-aware multi-load, multi-source coordination preserving comfort | Operational policy library callable by planning agents |

**Table 20.** Reported quantitative outcomes, baselines, system sizes and platforms of the examined studies for which the source reports them, with the evidence class of Table 21.

| Study | System and Size | Platform or Hardware | Baselines | Reported Outcome | Class |
|---|---|---|---|---|---|
| Diao et al. [27] | Converter-interfaced microgrid nodes; several simulated topologies; two-converter laboratory test | Simulation and a two-converter laboratory set-up | Coordination through an exogenous data channel | Coordination sustained with the exogenous channel removed, the remote information being inferred from the power exchanges by spiking networks trained online; the reported quantity is the elimination of the dedicated channel rather than a reduction ratio | E2 to E4 |
| Parisio et al. [92] | One experimental microgrid | Physical experimental microgrid | None; external feasibility study | Receding-horizon schedules feasible on the physical system and respecting the operating constraints; no numerical comparison is retained here because the economic comparison could not be confirmed from the source | E4 |
| Ioannou et al. [149] | Nano-grid with fluctuating generators, fluctuating loads and storage and one BDIx agent per asset | Simulation campaign over fluctuating photovoltaic generation and stochastic consumption variables | Centralized management | Sustained supply–demand balance with faster local reactions and lower coordination overhead than the centralized alternative | E2 |
| Ioannou et al. [177] | Aggregated single-bus mini-grid with four heterogeneous battery systems totaling 500 kWh | MATLAB R2025a simulation; three-day controller comparison and 24 h reward-weight sensitivity | Thirteen alternatives within a fourteen-controller comparison spanning rule-based, model-predictive and deep-reinforcement-learning families | Lowest state-of-charge variance of the fourteen, reported as essentially zero; lowest bus-voltage variance; near-minimal operating cost; moderate energy throughput; and robustness under high net-load volatility and communication impairments | E2 |
| Ioannou et al. [178] | Six-node proxy graph from the UCI household dataset with synthetic inconsistency episodes; simulated IEEE 13-node feeder | Offline evaluation; OpenDSS for the feeder | Eight baselines: LSTM, gated recurrent unit (GRU), spatio-temporal graph convolutional network (STGCN), diffusion convolutional recurrent neural network (DCRNN), Graph WaveNet, multivariate time-series graph neural network (MTGNN), Transformer and autoencoder | AUC-ROC of 0.8959, F1 score of 0.8307 and false-alarm rate of 0.41 percent on the proxy; F1 score 2.2 percent higher than GRU and up to 15.2 percent higher than LSTM, false-alarm rate about 52 percent lower than LSTM; AUC-ROC of 0.9016 and F1 score of 0.8361 on the 13-node feeder | E1 |
| Javaid et al. [147,171] | Nano-grid with fluctuating sources and loads | Physical implementation | None; external implementation study | Designated flow patterns maintained and voltage kept stable under fluctuation | E4 |
| Muhammad et al. [181] | Four interconnection scenarios among generators, loads and storage | Simulation with measured seasonal photovoltaic and fuel-cell data | Three control algorithms: total demand, adaptive demand and grid support | Storage-related losses reduced as interconnection richness increases; full daily demand met when grid support is available | E1 to E2 |

**Table 21.** Class position of every examined study on the unified scale of Table 6.

| Examined Study | Horizon | Authority | Evidence | Scale | Communication |
|---|---|---|---|---|---|
| *Control, optimization and multi-agent foundations (Table 14)* | | | | | |
| Lasseter [1] | T1 to T2 | n.a. | E0 | n.a. | n.a. |
| Guerrero et al. [4] | T1 to T3 | A5 | E2 to E4 | n.r. | C1 |
| Bidram and Davoudi [5] | T1 to T3 | n.a. | E0 | n.a. | n.a. |
| Parisio et al. [92] | T3 | A4 | E4 | n.r. | C1 |
| Cominesi et al. [93] | T2 to T3 | A3 | E2 | n.r. | C1 |
| Dimeas and Hatziargyriou [11] | T3 | A3 | E1 to E2 | n.r. | C1 |
| McArthur et al. [9,10] | n.a. | n.a. | E0 | n.a. | n.a. |
| Nunna and Doolla [12] | T3 | A3 | E2 | n.r. | C1 |
| Morstyn et al. [104] | T2 | A3 | E2 | n.r. | C1 |
| Zheng et al. [105] | T3 | A3 | E2 | n.r. | C2 |
| Howell et al. [68] | n.a. | n.a. | E0 | n.a. | C3 |
| *Reinforcement learning and coordination (Table 15)* | | | | | |
| Kuznetsova et al. [161] | T3 | A3 | E2 | n.r. | C1 |
| Foruzan et al. [119] | T3 | A3 | E2 | n.r. | C1 |
| Zeng et al. [120] | T3 | A3 | E2 | n.r. | C1 |
| Ji et al. [121] | T3 | A3 | E2 | n.r. | C1 |
| Du and Li [122] | T3 | A3 | E2 | S4 | C1 |
| Goh et al. [162] | T3 | A3 | E2 | n.r. | C1 |
| Ye et al. [123] | T3 | A3 | E2 | n.r. | C1 |
| *Semantic and task-oriented communication (Table 16)* | | | | | |
| Diao et al. [27] | T1 to T2 | A5 | E2 to E4 | S1 in the laboratory; n.r. in simulation | C2 |
| *Language-model and agentic systems (Table 17)* | | | | | |
| Grid-Agent [33] | T3 to T4 | A2 | E2 | S3 | C1 |
| X-GridAgent [34] | T4 | A1 | E1 | n.r. | C1 |
| Semantic VPP framework [167] | T3 | A2 | E1 to E2 | n.r. | C2 |
| Kiasari and Aly [35] | n.a. | n.a. | E0 | n.a. | n.a. |

**Table 21.** *Cont.*

| Examined Study | Horizon | Authority | Evidence | Scale | Communication |
|---|---|---|---|---|---|
| *Digital twins, neuro-symbolic assurance and federated foundation models (Table 18)* | | | | | |
| Islam et al. [142] | T2 to T3 | A3 | E3 to E4 | S2 | n.r. |
| Bruining et al. [60] | n.a. | n.a. | E0 | n.a. | n.a. |
| Addo et al. [62] | T2 to T3 | A3 | E2 | n.r. | C1 |
| Grid and federated foundation models [52,125] | T4 | A1 | E1 | n.r. | C2 |
| Alcántara and Chatzivasileiadis [54] | T4 | A1 | E1 | n.r. | C1 |
| *DAI/BDIx agent frameworks and flexible power-interface control (Table 19)* | | | | | |
| Ioannou et al. [149] | T2 to T3 | A3 | E2 | S2 | C1 |
| Ioannou et al. [176] | T3 to T4 | A1 | E1 to E2 | S2 | C1 |
| Ioannou et al. [177] | T2 to T3 | A3 | E2 | n.r. | C2 |
| Ioannou et al. [178] | T4 | A1 | E1 | n.r. | C1 |
| Javaid et al. [147,171] | T1 to T2 | A5 | E4 | S2 | C1 |
| Javaid et al. [148,179,180] | T2 to T3 | A2 | E1 | S2 | C1 |
| Muhammad et al. [181] | T3 | A3 | E1 to E2 | n.r. | C1 |

Rows in italics are group headings.

Taken together, the two lines occupy the two ends of the authority ladder developed later in Section 4.7.1. Deliberation, forecasting and negotiation are performed by the DAI/BDIx agents at management horizons, whereas execution and verification are performed by the flexible power-interface methods at device horizons. Complementary approaches to distributed power management, predictive-adaptive control and power-flow assignment are developed in these studies [149,177,181]; their combination with semantic contracts and proof-carrying execution remains to be evaluated. As in the other clusters, the semantic content, acknowledgment state and adversarial robustness of the exchanged agent messages have not yet been formalized. These limitations motivate the semantic-control and assurance layers of the proposed architecture. The outcomes on which the class positions of the two lines rest are collected, together with the baselines against which they were obtained, in Table 20.

Table 19 compares two research lines that are complementary in both abstraction and time horizon. The DAI/BDIx studies of Ioannou et al. concentrate on the organization of distributed decision making. Their recurring contribution is to place an autonomous reasoning cycle at each participating device or energy asset, allowing beliefs, objectives and selected plans to be maintained locally while coordination emerges through agent interaction [149]. The later studies enrich this organization with agent-embedded forecasting [176], federated forecasting with adaptive correction [177] and physics-guided graph forecasting of electricity inconsistencies [178]. Across the line, distributed organization, forecasting and adaptive control are investigated in different experimental settings; the reported outcomes apply to the respective modeled systems and evaluated baselines. The remaining limitation is that agent messages, commitments and authority are not yet represented through a standardized semantic contract with explicit common-knowledge and adversarial-failure tests.

The Javaid et al. line begins closer to the electrical interface. Power-flow coloring makes individual source-to-load flows explicit under fluctuating sources and loads [171], while the cooperative distributed control implementation demonstrates how designated flow patterns can be maintained over a nano-grid with fluctuating loads [147]. The balanceability studies then derive analytical conditions under which a configuration can safely absorb specified fluctuation ranges [148], reduce the effort of checking those conditions [179] and determine the storage required to guarantee the property [180]. The later loss-mitigation work adds richer interconnection patterns and critical-load priorities [181]. Relative to the DAI/BDIx studies, this line offers less flexible deliberation but stronger device-level interpretability and checkability.

The principal comparison therefore concerns how the two research lines can be composed. BDIx agents can interpret objectives, maintain local beliefs, negotiate commitments and select a candidate plan. Flexible power-interface and balanceability methods can translate that plan into accountable flows and determine whether the physical configuration can execute it safely. The analytical conditions can become predicates in a proof-carrying certificate, while the agent organization supplies the distributed reasoning that constructs the proposal. A complete evaluation should combine both lines in the same closed-loop scenarios and report management quality, message overhead, semantic agreement, balanceability margin, transient response, certificate latency and behavior under communication loss or compromised agents. Such an experiment would directly test the proposed bridge between deliberative distributed intelligence and certified electrical execution.

### *4.7. Cross-Study Comparison and Evidence Quality*

The cluster examinations are completed by comparing control horizon, information assumptions, assurance evidence and the resulting evidence gap across the research fam-

ilies. A method may appear mature within machine learning because it outperforms a baseline yet remain immature for power-system deployment because the electrical model is static, communication is ideal, safety is expressed only through penalties and no real-time implementation is tested. Conversely, a classical distributed controller may have strong stability evidence but provide little support for semantic context, adaptation or human interaction. Complementary evidence is therefore identified rather than one universally superior family being selected.

#### 4.7.1. Control Horizon, Authority and Physical Validation

The first distinction is where a method sits in the control hierarchy. Primary and inner converter controls operate at sub-second or millisecond time scales and require deterministic execution. Secondary control restores voltage and frequency over slower but still tightly bounded intervals. Tertiary and EMS decisions range from seconds to hours. Most RL energy-management studies belong to the latter category, even when described as "real time", because their actions are scheduling or set-point decisions rather than pulse-width-modulation or protection actions. LLM-based systems are slower still and are best treated as supervisory planning and analysis components. The horizon, authority and validation position of each research family is summarized in Table 22. The same table carries the evidence class normally reached by the family, the evidence that is most frequently absent from it, the items without which its results cannot be reproduced and the minimum claim that its evidence supports so that a family is characterized once and in one place rather than described separately under each comparison axis. The classes reported in the table are those of Table 6; the entries summarize the authority and evidence range represented by the selected studies and do not imply that every study reaches the highest class. The class position of every examined study is recorded first in Table 21 and the family-level entries of Table 22 are read from it.

The conventions of Table 21 are described as follows. The authority class records the authority actually exercised in the reported evaluation and not the authority intended by the design, so a method whose actions were applied only inside a simulator is recorded at A3 even when its design targets set-point application. Evidence classes describe the reported validation environments. A range records multiple environments within one publication or across grouped publications; it does not imply that every intermediate evidence level was tested. The n.r. entry means that the available extraction does not establish the item in a form that permits confident assignment; it is not proof that the original publication omitted it. An n.a. entry means that the class does not apply because the work is a survey, definition or architectural proposal. The table contains 35 grouped entries: six classified at E0 and 29 with an evaluation entry. Scale is entirely unassigned in 21 of those 29 entries and partly unassigned for the simulation part of the Diao entry. These are extraction limitations rather than verified omissions from 21 publications. The seven reinforcement-learning entries are classified as C1 in this extraction and E3 is not established for the selected reinforcement-learning, agentic or foundation-model entries. The corresponding family ranges are summarized in Table 22.

Validation maturity should be interpreted relative to this horizon. An optimization algorithm tested on day-ahead profiles provides useful economic evidence, but it cannot support a claim about transient stability. A semantic codec evaluated through task accuracy provides communication evidence, but it cannot support a claim about safe load shedding unless the receiving controller and physical transients are in the loop. The strongest future studies will therefore combine several simulators or test layers: an energy-management environment, a dynamic or electromagnetic-transient model, a communication network and an agent/tool environment.

**Table 22.** Unified cross-family evaluation matrix, stated in the classes of Table 6.

| Research Family | Horizon and Typical Output | Highest Defensible Authority, Conditional on the Evidence Class | Evidence and Common Validation | Frequently Missing Evidence | Reproducibility Essentials | Minimum Defensible Claim and Architectural Role |
|---|---|---|---|---|---|---|
| Droop and hierarchical control [4,5] | T1 to T2; converter references, restoration signals and power commands | A5, direct actuation within the designed operating envelope | E2 to E4; analysis, dynamic simulation and experiments | Semantic mismatch and malicious communication for connected implementations; broader asset and market uncertainty in some studies; artificial-intelligence governance absent or not applicable | Plant parameters, gains, topology and disturbance scripts | Stability and performance within the stated model and communication assumptions; edge-reflex and verified control foundation |
| Distributed MPC and ADMM [92,93,105] | T2 to T3; constrained set points and schedules | A3 in simulation to A4 in the selected physical implementation | E2 to E4; optimization simulation and selected physical experiments | Fast dynamics and protection effects; model mismatch, delays, missing values, semantic provenance and fallback testing; solver and model failure handling | Solver, horizon, forecasts, scenarios and infeasibility policy | Feasible and economic schedules for the evaluated scenarios; numerical planning tool for distributed agents |
| Reinforcement learning for energy management [121,123,161] | T2 to T3; charge, discharge, dispatch, curtailment and market actions | A3, actions applied in the tested simulator; physical authority is not established | E2; closed-loop simulation in the selected studies | Hardware-in-the-loop and rare-event tests; safe exploration and distribution-shift evaluation; asynchronous and partial observations; formal safe-set coverage and update governance | Environment, rewards, seeds, networks, training curves and failure cases | Performance in the tested simulator and not deployment safety; adaptive policy proposer |
| Multi-agent reinforcement learning and common-knowledge methods [15,134] | Task dependent; decentralized coordinated actions | A3, simulation-level action unless communication and safe-set assumptions are verified | E0 to E2; generic multi-agent benchmarks or power simulations | Grid-specific communication, stability and authority validation; realistic partial observability | Environment, rewards, seeds, networks, training curves and failure cases | Coordination performance in the evaluated benchmark; distributed coordination and belief-aware planning |
| Semantic communication [23,27] | T1 to T3; compressed features, semantic states or intents | Information outputs alone carry no plant authority; the selected converter implementation includes A5 | E1 to E4 across communication studies and the selected physical grid case | Closed-loop electrical consequences; ontology drift, source authority and common knowledge; safety preservation after compression | Dataset, channel, encoder, task loss, codebook and mismatch tests | Task utility and bandwidth result under the evaluated channel and task; semantic and common-knowledge layer |

**Table 22.** *Cont.*

| Research Family | Horizon and Typical Output | Highest Defensible Authority, Conditional on the Evidence Class | Evidence and Common Validation | Frequently Missing Evidence | Reproducibility Essentials | Minimum Defensible Claim and Architectural Role |
|---|---|---|---|---|---|---|
| Language-model agents [33,34] | T3 to T4; plans, tool calls, explanations and candidate corrective sequences | A2, advisory or sandboxed proposal | E0 to E2; standard network simulations and query tasks | Timing, protection and closed-loop field behavior; typed interfaces, cyber isolation, independent verification and operator studies; stale retrieval, prompt attacks and cross-agent trust; formal authority, verifier coverage and rollback | Model version, prompts, tools, schemas, retrieved corpora and traces | Assistance or sandboxed planning performance in the evaluated cases; agentic and twin supervisory layer |
| Digital twins and neuro-symbolic assurance [61,62,142] | Model dependent; predictions, counterfactuals, rule decisions and certificates | A2 to A3, validation and advisory unless the twin is synchronized and qualified | E0 to E4 across conceptual treatments, simulations and the selected laboratory twins | Fidelity bounds, synchronization, failover and closed-loop qualification; pilot evidence; data provenance and twin-state consistency; rule completeness and verifier independence | Models, parameter updates, uncertainty, rules and synchronization logs | Decision support within the calibrated twin scope; assurance and consequence-prediction layer |
| Federated and foundation models [17,52] | T4; shared representations, forecasts or analytical outputs | A1, decision support | E0 to E1; offline datasets and simulations | Closed-loop effect of model errors; site heterogeneity and update communication; continual monitoring, poisoning defense, model registry and field adaptation; privacy and lifecycle control | Client partitions, rounds, participation, privacy settings, attacks and checkpoints | Predictive or analytical performance under the reported federation; distributed learning infrastructure |
| DAI and BDIx agent frameworks [149,176–178] | T2 to T4 across management and analytical tasks, with sub-second reflexes retained by classical control; per-asset decisions, forecasts and bounded balancing set points | A1 for analytical forecasts to A3 for simulated actions; no physical authority established | E1 to E2; analytical evaluations and power simulations with study-specific centralized, rule-based, model-predictive and learning baselines | Hardware-in-the-loop, semantic-failure and adversarial-communication testing; heterogeneous field conditions | Agent roles, belief and plan libraries, baselines, network and load data and seeds | Management-horizon coordination against the reported baselines; distributed-intelligence and agentic coordination layers |
| Flexible power interface and balanceability control [147,148,171,179–181] | T1 to T2; per-flow power assignments and checkable balanceability conditions | A5, direct actuation of the interface converters within proven balanceable sets | E1 to E4; analytical conditions with numerical checks, nano-grid implementation and measured-data studies | Cyber, semantic-mismatch and multi-vendor interoperability tests | Interface topology, conversion limits, balanceability predicates and measured profiles | Feasibility and balanceability within the proven operating set; verified execution layer and certificate predicates for assurance |

The reviewed literature also differs in action authority. Controller outputs are typically defined precisely in classical control papers. Set-points or schedules are commonly produced in RL studies, while multi-step commands or tool requests may be generated in language-agent papers. A stronger validator is required as the action space becomes more expressive. An authority ladder is therefore imposed by the proposed architecture and it is the A1 to A5 ladder of Table 6. The highest evaluated level should be stated rather than the broad "autonomous" label being used. The two approach families examined in Section 4.6 were placed deliberately at complementary rungs of this ladder. Agent deliberation in the DAI/BDIx frameworks is confined to management horizons and bounded set points [149]. Deterministic device-level authority backed by checkable feasibility conditions is retained by power-flow assignment [171] and balanceability checking [148].

Table 22 demonstrates that algorithm families should be compared only after their physical horizon and action authority have been aligned. Primary and protection functions operate at sub-cycle to sub-second scales and directly influence voltage, current, frequency and switching. Their evidence must therefore include real-time execution, transient models and preferably HIL or power-HIL. Scheduling, market and asset-management methods act over minutes to hours and can be evaluated initially through optimization or closed-loop simulation, although they still require a defined fallback when forecasts or solvers fail. Semantic communication and agentic planning span several horizons because the same technology can be used either to summarize a slow market request or to transmit a rapid event. The permissible authority must consequently be specified per use case rather than inferred from the technology name.

The table also exposes why impressive planning accuracy cannot justify direct control authority. LLM agents and foundation models operate primarily at supervisory or analytical horizons and their most defensible role is the production of advice, structured candidate plans or tool calls. Digital twins and neuro-symbolic systems can move closer to control only when synchronization error, model discrepancy, rule coverage and verifier independence are quantified. DAI/BDIx frameworks occupy an intermediate management layer, whereas flexible power-interface methods can act closer to the plant because their commands and feasibility conditions are more explicit. A fair comparison should therefore hold authority constant: candidate planners should be compared with candidate planners, certified schedulers with certified schedulers and fast controllers with fast controllers. This prevents a slow high-level system from appearing superior simply because the evaluation omits the transient and protection responsibilities borne by a lower-level controller.

#### 4.7.2. Communication, Semantics and Knowledge Assumptions

The second distinction concerns what each family assumes about communication and knowledge. Classical distributed control usually specifies topology, delays or connectivity mathematically, but it treats the exchanged value as unambiguous. Semantic communication studies analyze meaning and task utility, but they often abstract the physical consequences of a misunderstood message. LLM-agent systems reason over natural language and tool schemas, yet they may assume that retrieved data are current and that the selected tool corresponds to the user's actual objective. The cognitive mini-grid sits at the intersection of these assumptions and therefore needs an explicit message contract.

The minimum contract contains a source identifier, authenticated role, message type, semantic version, physical scope, timestamp, validity horizon, confidence, units, provenance, requested acknowledgment and commitment status. A statement such as "battery reserve available" is not actionable until the receiving agent knows which battery, how much energy and power are available, for how long, under which state-of-health and

emergency-reserve policy and whether the sender is authorized to commit it. This requirement connects ontologies with control engineering.

Four epistemic states should be distinguished. Local knowledge is available to one agent. Shared information has been transmitted to a group but may not have been received or interpreted by every member. Mutual knowledge means the relevant agents know the fact. Operational common knowledge means the agents know the fact and possess sufficient evidence that the others know it for the purpose and validity interval of a specified coordinated action. The last state normally requires acknowledgments or a trusted replicated record; it should not be inferred from an unaudited broadcast. The condition is made testable below, with its communication and failure assumptions, its commit protocol and a coordinated switching example.

The operational condition is specified for a proposed action ($a$), participant set ($P$), state description ($s$), ontology version ($v$), execution time ($t_{\text{exec}}$) and requested interval ($[t_{\text{exec}}, t_{\text{exec}} + \tau]$). A designated committer requires a signed preparation acknowledgment from every member of $P$ before the deadline ($t_c < t_{\text{exec}}$). The signed tuple binds the plan identifier, participant set, step identifiers and shared state and timing fields; persistent duplicate rejection prevents replay. A commit record containing those acknowledgments is written to a protected committer log and distributed. Replication, if used, requires its own consistency, fault and recovery model. Holding the record establishes auditable preparation readiness for a member; it does not establish that every other member received the commit or will execute it. Receipt acknowledgments provide delivery evidence to the committer. A missing receipt acknowledgment leaves delivery unknown, so the committer must consider both possible outcomes. Halpern and Moses explain why exact common knowledge cannot, in general, be obtained by a finite acknowledgment chain when communication is not guaranteed [115]; their result does not prove the safety of the switching protocol proposed here.

Messages may be lost or delayed without bound, agents may crash at any time and clocks are assumed to remain within a specified skew ($\delta$). Delivered records must pass authentication, integrity, authorization and replay checks. Independent trusted Layer 2 measurements are required for the physical prerequisites of each step. A member without a valid commit record does not execute, but partial delivery or a crash can still cause partial execution. Every dependent step must therefore be gated on verified predecessor effects, local protection conditions and a certified admissible intermediate state. Zero current alone does not prove an open switch. Trusted breaker-position, isolation and appropriate synchronism or dead-bus checks are required. The protocol in Table 23 specifies these design requirements; it does not establish atomic execution or unconditional safety under the stated failures. Plans whose safety requires simultaneous actions that cannot tolerate partial execution must not be authorized through this protocol.

The example of Table 23 answers the question of which agents must know that the action has been committed: every member of $P$ whose step is irreversible or whose step makes the step of another member safe must hold the commit record and the members whose steps are gated on physical confirmation need only the record and the confirmation, not the knowledge of the other members' knowledge. In the example, SW11 holds the commit record although its receipt acknowledgment was lost, so it may execute its gated step; the committer, lacking that acknowledgment, records that the delivery to SW11 is unknown and prepares the recovery decision for both outcomes. The protocol does not establish atomic execution. The effect of a missing acknowledgment must instead be handled by the independently verified physical gates and recovery policies.

**Table 23.** Proposed preparation and commit protocol for a planned feeder transfer, with partial execution handled through independently verified local gates.

| Phase | Rule | Coordinated Switching Example: Transfer of Zone Z5 from Feeder A to Feeder B by Opening SW9 and Closing SW11 |
|---|---|---|
| Proposal | The committer sends $(s, a, v, t_{\text{exec}}, \tau)$ to every member of $P$ with the commit deadline ($t_c$) | The substation agent sends the transfer plan, ontology version 4, $t_{\text{exec}}$ at 12:08:00 and $\tau$ of 5 min to the agents of SW9, SW11 and storage B4, with $t_c$ at 12:07:30 |
| Acknowledgment | Each member checks the ontology version, its own state and its authority and returns a signed acknowledgment of the identical tuple; a member that cannot acknowledge returns a signed refusal or nothing | SW9 confirms it is closed and operable, SW11 confirms it is open and operable, B4 confirms 150 kW of headroom for 5 min; all three sign the same tuple |
| Commit | If every acknowledgment is held before $t_c$, the committer writes the signed commit record with all acknowledgments to the protected committer log and sends it to every member; otherwise, it writes an abort record | All three acknowledgments arrive by 12:07:20; the commit record is written and sent |
| Confirmation | Each member acknowledges receipt of the commit record; the acknowledgment confirms receipt but does not exclude a later crash. A member without a valid record at $t_{\text{exec}}$ does not execute | SW9 and B4 acknowledge; the acknowledgment of SW11 is lost |
| Execution | At $t_{\text{exec}}$, a member holding the commit record executes its step; a step whose safety depends on a preceding step waits for the physical confirmation of that step at Layer 2 within a bounded window and otherwise does not execute | SW9 opens at 12:08:00 if its local checks pass and B4 dispatches only under its certified conditions. SW11 may close within the 2 s confirmation window only after trusted breaker-position and isolation checks establish the required SW9 state and synchronism or dead-bus checks pass; otherwise, SW11 remains open |
| Failure case | A member that did not receive the commit record does nothing; a member whose confirmation window expires does nothing; partial execution is possible, so the resulting measured state must satisfy a certified hold condition or a separately verified recovery policy | If SW11 receives no valid commit, it remains open. A state with both switches open may be held only after isolation and local supply or reserve conditions are verified; reclosing SW9 requires a fresh protection, topology and synchronization check |

The assumptions in Table 24 form a minimum semantic message contract. The table groups the twelve message fields stated above into seven assumption categories rather than reproducing one row for every field. Timeliness and validity determine whether the described state still supports the intended action. Semantic equivalence determines whether all recipients use the same units, topology version, reserve definition and operating policy. Delivery and acknowledgment determine whether a coordinated action can be considered mutually understood rather than merely broadcast. Source authority determines whether an agent is permitted to commit the asset or issue the request. Confidence and provenance distinguish direct measurement from forecast, inference or operator statement. Privacy limits the amount of operational and behavioral detail that may be disclosed. Common knowledge then combines these properties for a particular group action and time interval.

**Table 24.** Information assumptions that must be reported in cognitive mini-grid studies.

| Assumption | Common Implicit Form | Why It Matters Physically | Required Experiment/ Reporting Item |
|---|---|---|---|
| Timeliness | All observations represent the current state | A stale state may reverse the sign or feasibility of a corrective action | Timestamp, age-of-information distribution and deadline-miss behavior |
| Semantic equivalence | All agents assign the same meaning to a code/message | Different units, reserve definitions or topology versions can produce conflicting actions | Ontology/codebook version, mismatch injection and fallback response |
| Delivery and acknowledgement | Broadcast implies shared knowledge | Partial delivery can make coordinated switching or dispatch unsafe | Packet loss, partition, acknowledgement protocol and abort rule |
| Source authority | A valid message is authorized | A compromised or mis-scoped agent can commit unavailable assets | Role credentials, least privilege and rejected-command statistics |
| Confidence and provenance | Model output is treated as a measurement | Forecasts, inferred states and operator statements carry different uncertainty | Source type, uncertainty, calibration and provenance chain |
| Privacy | Shared features reveal no sensitive information | Load patterns, bids and maintenance status can expose users or vulnerabilities | Threat model, privacy mechanism and utility/privacy trade-off |
| Common knowledge | Agents coordinate from a public signal | Different observations or filters may create inconsistent beliefs | Definition of operational common knowledge and explicit establishment test |

These assumptions are tightly coupled. A message can be authentic but stale, semantically correct but unauthorized, timely but delivered to only part of the required coalition or widely shared but derived from an uncalibrated model. Testing one property in isolation is therefore insufficient. Experiments should inject compound failures such as delayed but correctly signed messages, ontology mismatch combined with partial acknowledgment and high-confidence forecasts generated from shifted data. The controller should report whether it proceeds, requests additional evidence, falls back to raw telemetry, reduces authority or aborts the coordinated action. The comparison also implies that "common knowledge" cannot be claimed from a shared broadcast alone. The establishment test succeeds only when authenticated identities, the ontology and its version, the relevant state, the validity interval, the required acknowledgments and the abort rule have all been recorded. If a critical preparation acknowledgment is missing or inconsistent, the action is not committed. A missing receipt acknowledgment leaves commit delivery unknown and must be handled under the partial-execution and recovery rules.

4.7.3. Safety, Explainability and Cybersecurity Evidence

The third distinction concerns the assurance evidence that each family provides. The reviewed fields use different meanings of safety. In optimization, safety often means constraint feasibility for the modeled horizon. In RL, it may mean a low rate of violations or a constrained return. In formal methods, it means a proven invariant under stated assumptions. In cybersecurity, it means resistance, detection and recovery under an adversary. In operator studies, it includes understandable authority and effective override. A cognitive mini-grid paper should state which meaning is evaluated and avoid substituting one for another.

Explainability is likewise multi-level. Feature importance can explain a prediction, but it does not explain why a sequence of tool calls was selected. A natural-language rationale can be readable but unfaithful to the actual computation. A stronger explanation package contains the triggering evidence, retrieved sources, tool inputs and outputs, checked constraints, rejected alternatives and the final authorization decision. This package also supports incident reconstruction and regulatory accountability.

Cybersecurity must cover both established operational-technology threats and agent-specific threats, a requirement developed further in Section 6.5. The former include false-data injection, denial of service and unauthorized control, together with supply-chain compromise. The first three are cataloged in the cyber-physical security surveys of the smart grid [182,183] and of cyber-physical systems in general [184], in the secure-control agenda for survivable cyber-physical systems [185] and in the vulnerability assessment of supervisory control and data acquisition (SCADA) systems [186]. The latter include prompt injection, poisoned retrieval corpora, malicious tool descriptions, cross-agent impersonation, memory manipulation and adversarial semantic codes. Existing guidance remains necessary. The IEC 62351 security standards [151]; the National Institute of Standards and Technology Interagency Report (NISTIR) 7628 smart-grid cybersecurity guidelines [152]; and the NIST Special Publication (SP) 800-82 guide to operational-technology security, which also addresses supply-chain risk [158], form this baseline. Model provenance, prompt/tool isolation and signed action certificates nevertheless require additional controls.

The final three columns of Table 22 are intended to discipline the claims that can be made from each evidence family and each family is strong in a different place: explicit models and stability analysis in classical and distributed control, explicit constraints and an interpretable objective in optimization and predictive control, adaptation in high-dimensional spaces in reinforcement learning, communication and task co-design in semantic communication, flexible planning and tool orchestration in language-model agents, counterfactual analysis in twins and neuro-symbolic methods, cross-site representation reuse in federated and foundation models, bounded management-horizon coordination in the agent frameworks and checkable device-level feasibility in the flexible power-interface methods. Classical and distributed-control studies can support strong statements about stability and performance within their modeled plant and communication assumptions, but they do not automatically establish robustness to semantic ambiguity or malicious communication. Optimization and MPC studies can support claims about feasibility and economic performance over the evaluated horizon, provided that solver settings, forecast scenarios and infeasibility handling are reported. RL and MARL studies normally support performance claims within the tested environment; without rare-event, hardware and safe-set evidence, they should not be generalized to deployment safety. Semantic-communication studies support task-utility and communication-efficiency claims under their evaluated task and channel, not necessarily preservation of electrical safety. LLM-agent studies support assistance or sandbox-planning claims unless authority, timing and independent validation are demonstrated.

The matrix also shows that reproducibility requirements are family-specific. Publishing code without plant parameters, disturbance scripts or controller gains is insufficient for control studies. Reporting an LLM name without model version, prompts, tool schemas, retrieved documents and traces is insufficient for agentic studies. A federated-learning result cannot be reproduced without client partitions, participation patterns, aggregation rounds, privacy settings and attack assumptions. A digital twin requires synchronization logs, update rules and uncertainty bounds, in addition to the model itself. The minimum defensible claim in the final column should therefore appear explicitly in future manuscripts. Stronger claims should be accompanied by the additional evidence identified in the missing-evidence columns, allowing readers to distinguish a useful conceptual prototype from a deployment-grade contribution.

#### 4.7.4. Critical Appraisal of Methodological Limitations and Applicability Boundaries

The comparison so far records what each family establishes. The appraisal that follows records why several of the reported results do not extend as far as their phrasing suggests and it is stated at the level of the individual study wherever the limitation originates in a specific design choice rather than in the family as a whole. The family-level summary is given in Table 25.

In the control and optimization literature, the recurring limitation is not the analysis but the scope of the plant model to which the analysis is attached. The hierarchical control results of Guerrero et al. [4] and the secondary-control results derived from them are proven for a converter model with specified impedances and a fixed topology, so their guarantees are conditional on that model and do not, by themselves, cover reconfiguration, unbalanced operation or the loss of a communication link that the same papers do not model. The predictive-control experiment of Parisio et al. [92] establishes that the scheme is feasible on a physical microgrid and that its schedules respect the operating constraints, which is a stronger result than a simulation, but the experiment is conducted on one installation and the economic behavior under a different tariff, asset mix or forecast quality is outside what the experiment can support. The multi-agent controller of Dimeas and Hatziargyriou [11] demonstrates that local dispatch can be settled without a central decision process, although the demonstration is a market mechanism evaluated on a small configuration and the resilience claims that are frequently attributed to multi-agent operation are not tested in it.

In the learning literature, the recurring limitation is that the environment in which the policy is trained is also the environment in which it is judged. The energy-management study of Kuznetsova et al. [161] reports its outcome across the simulated scenarios that were generated for it, so the result characterizes the policy within that generator and not against an independently specified baseline. The safety-constrained formulation of Ye et al. [123] reduces constraint violations substantially, but a reduction in violations is a statement about the frequency of an event in a simulator and not a proof of an invariant and it therefore does not transfer to a claim that the constraint is never violated on a plant. The improvement reported by Lee et al. [163] is stated against the baselines that were compared in that study, which reflects a common pattern within this family. The strength of a learned policy is normally established against baselines selected and configured within the same work. The margin therefore characterizes that comparison rather than the method in all settings. None of the reinforcement-learning studies examined here reports an evaluation under communication degradation, so their conclusions are bounded to the class C1 and C2 assumptions of Table 6.

**Table 25.** Principal methodological limitations, experimental limitations and applicability boundaries by research family.

| Research Family | Principal Methodological Limitation | Experimental Limitation | Boundary of Applicability of the Conclusions |
|---|---|---|---|
| Droop and hierarchical control | Guarantees are conditional on a fixed converter model, impedance set and topology | Experiments are conducted on one installation and one configuration | Valid within the modeled envelope; not extended to reconfiguration, unbalance or link loss |
| Distributed MPC and ADMM | Optimality and feasibility are asserted for the modeled horizon and the assumed forecast quality | Solver failure, infeasibility handling and forecast degradation are seldom exercised | Valid for the evaluated scenario set and tariff structure |
| Reinforcement learning for energy management | The training environment is also the evaluation environment and baselines are tuned by the same authors | Rare contingencies, transient dynamics and hardware are absent | Valid inside the simulator and under ideal or modeled communication |
| Multi-agent reinforcement learning and common-knowledge methods | Coordination gains are demonstrated on abstract benchmarks whose observation structure differs from a feeder | Power-system communication, protection and authority constraints are not represented | Valid as evidence that a coordination mechanism exists, not that it is admissible in a grid |
| Semantic communication | Task utility is optimized without a model of the electrical consequence of a misinterpreted message | Validation is at device or link scale, with few concurrent senders | Valid for the evaluated channel, task and encoder pair, not for feeder-scale semantic consistency |
| Language-model agents | Success rates are computed on author-constructed task sets rather than an independently specified population | Timing, protection interaction and closed-loop behavior are not measured | Valid as feasibility evidence for sandboxed planning in the evaluated cases |
| Digital twins and neuro-symbolic assurance | No accepted fidelity criterion exists and the verifier is frequently derived from the model it verifies | Live twins are demonstrated at laboratory scale; rule coverage is not established | Valid within the calibrated scope of the twin and the encoded rule set |
| Federated and foundation models | Aggregate accuracy is reported without the closed-loop consequence of a model error | Site heterogeneity, participation patterns and poisoning are rarely evaluated together | Valid as predictive performance under the reported federation only |
| DAI and BDIx agent frameworks | Comparison is made against reimplemented centralized, droop and predictive baselines | Hardware-in-the-loop and adversarial-communication evidence is not yet reported | Valid at the management horizon and the simulated system scale |
| Flexible power interface and balanceability control | Feasibility conditions are proven for the modeled interface and conversion limits | Multi-vendor interoperability and cyber degradation are not exercised | Valid within the proven balanceable set of the modeled interface |

In the semantic and agentic literature, the limitation is one of scale and of closure. The neuromorphic semantic scheme of Diao et al. [27] is validated in simulation across topologies and additionally in a two-converter laboratory test, which is genuine physical evidence, but a two-converter test cannot establish the behavior of the scheme at feeder scale, where the number of concurrent senders and the diversity of encoders are the quantities that determine whether the semantic representation remains consistent. The agentic prototypes [33,34] are evaluated on network simulations and query tasks and their reported success rate is a property of the task set that the authors constructed; because the tasks are not drawn from an independently specified population, the rate cannot be read as a reliability figure. The review of Kiasari and Aly [35] consolidates this literature and applies an evaluation framework to it, yet the underlying corpus that it consolidates is itself largely at evidence levels E0 to E2, so the maturity of the review does not raise the maturity of the field that it reviews.

In the twin and assurance literature, the limitation is the absence of an accepted fidelity criterion. The two laboratory-scale live twins reported by Islam et al. [142] demonstrate that synchronization is achievable at that scale and identify fidelity standards, multi-physics modeling and real-time synchronization as the open challenges, which is precisely why the results cannot be extrapolated: the quantity that would license extrapolation is the one the authors identify as undefined. The neuro-symbolic controller of Addo et al. [62] proposes control actions from a trainable actor–critic policy and reports improved performance in simulations on a distribution-grid testbed, so its conclusions are bounded to that testbed and to the rule set that was encoded and the completeness of that rule set is not established. Across the family, the verifier is frequently constructed by the same authors and from the same model as the proposer, which weakens the independence on which the assurance argument depends.

Two boundaries apply to the appraisal itself. The first is that an absent evaluation is not a defect when the contribution is conceptual: a paper that proposes a representation is not obliged to provide hardware evidence and the appraisal therefore records what a result supports rather than assigning a deficiency. The second is that the appraisal is conducted from the published record, so a limitation that the authors addressed without reporting it cannot be detected here. Both boundaries are reasons for the reporting requirements of Section 7.6 rather than qualifications of them.

#### 4.7.5. Synthesis of the Evidence Gap

The comparison reveals a consistent fragmentation, which confirms the gaps identified in Section 2.3. Power-control papers offer the strongest physical foundations but rarely model semantic meaning or agent deliberation. Semantic-communication papers optimize relevance and bandwidth but rarely carry the analysis through protection, stability and multi-vendor interoperability. Learning papers provide adaptation but usually depend on simplified simulators and reward-based safety. Agentic papers offer flexible workflow automation but remain early and primarily simulation-based. Digital-twin and neuro-symbolic research offers a pathway to validation but lacks widely accepted fidelity and certification procedures.

No examined stream, by itself, supplies all six layers of the architecture defined in Section 3.5 and none satisfies more than two of the interface contracts specified in Section 3.5.1; the limitations appraised in Section 4.7.4 are therefore complementary rather than shared. The cognitive power mini-grid should therefore be treated as a compositional research program. Progress requires interfaces that allow mature control, optimization and protection methods to remain authoritative while semantic, learning and agentic components contribute context, forecasts, candidate plans and explanations. The decisive

research question is not which AI model is best in isolation but whether the complete composition preserves stability and authority under uncertainty, communication failures, malicious inputs and model errors.

## 5. Applications

In this section, the proposed cognitive architecture is translated into five operational application classes. Namely, real-time balancing and adaptive energy management; resilience, islanding, black start and self-healing; semantic demand response and human preference translation; predictive maintenance and multimodal asset intelligence; and community energy markets and inter-mini-grid federations. For every application, the physical objective, required cognitive functions, information exchanges, validation needs and principal failure modes are examined, after which the relative maturity, assurance requirements and appropriate deployment order of the applications are compared. Through this organization, the operational value of semantic and agentic intelligence is assessed, together with the risks created when decisions span several assets, stakeholders or control layers.

### *5.1. Real-Time Balancing and Adaptive Energy Management*

Islanded mini-grids must balance supply and demand with limited inertia and reserve. Conventional optimization, represented by the MPC and distributed-optimization studies examined in Section 4.1, performs well when forecasts and models are reliable, while RL can adapt to uncertain prices, renewable generation and demand. The learning studies examined in Section 4.2 cover battery scheduling by tabular Q-learning [161], distributed multi-agent Q-learning for economic management [119], approximate dynamic programming with recurrent value approximation [120] and real-time scheduling by deep Q-learning [121]. They also cover multi-microgrid coordination [122], safe DRL [123], reward-design assessment [162] and an integrated DRL-based EMS architecture [163].

Cognitive balancing adds semantic prioritization and explicit coordination. Rather than sharing every device trajectory, agents can announce reserve states, ramping capability, confidence and binding constraints. A battery agent may publish "20 kWh dispatchable until 17:30, excluding emergency reserve", while a load agent may publish "10 kW deferrable for 30 min" and a coordinator can construct an action while preserving these meanings. Common knowledge is especially important when several agents must change power simultaneously to avoid rebound or oscillation.

The preferred architecture is dual-rate. Fast local controllers maintain electrical stability. A slower distributed optimization or MARL layer adjusts set points. An agentic layer interprets exceptional goals (for example, preserving a clinic load during a fuel shortage) and queries optimization tools. The result is a structured decision rather than free-form control.

### *5.2. Resilience, Islanding, Black Start and Self-Healing*

Resilience applications include disturbance detection, intentional islanding, fault isolation, service restoration, black start and reconnection. MASs are attractive because local agents can continue to operate after loss of a central controller. Control strategies for islanded operation were defined early [187], autonomous microgrid control without a communication network was demonstrated by Piagi and Lasseter [188] and the stability aspects specific to microgrids were analyzed by Majumder [189]. Distributed secondary control can restore frequency and voltage, whether through distributed averaging [76], the distributed approach of Shafiee et al. [190] or feedback-linearization-based cooperative control [191], but communication failures complicate the assumptions required for coordination, as the treatment of uncertain communication links by Lu et al. shows [192].

Resilience is bounded by the physical resources that exist before the disturbance and the cognitive layers cannot restore a load for which no energy, no path and no switching capability were provided. The restoration capacity of a mini-grid is therefore fixed at planning time by the energy and power ratings of its stationary storage; the black-start capability of its generators, the placement of the sectionalizing switches and soft open points that determine which sections can be re-energized; the reserve policy that decides how much stored energy is withheld for contingencies; and the availability of transportable resources. Namely, mobile energy storage, mobile generators and electric-vehicle fleets, whose siting, pre-positioning and routing are planning decisions rather than control decisions. The joint planning of transportable and flexible resources for renewable integration and resilience enhancement addresses precisely this coupling between the resources that are installed or pre-positioned and the resilience that operation can then deliver [193] and the framework for resilient community microgrids of Billanes et al. [69] consolidates the operational strategies and the performance indicators, including reliability, stability and flexibility measures, against which that resilience is assessed. In the architecture proposed here, these results enter at two points. At planning time, Layer 5 hosts the twin and the optimization tools through which the resource portfolio is sized, the resilience indicators are evaluated before commissioning and the reserve policy that results becomes part of the constraint set ($\mathcal{C}$) of the governor. At operating time, the transportable resources are represented as assets of Layer 1 with agents of their own in Layer 4 so that a mobile storage unit en route to a bus is the source of a belief with a validity horizon rather than an assumed reserve, as the restoration sequence of Section 3.5.3 shows. The resilience indicators consolidated by Billanes et al. are consistent with the metric families of Section 7.6.2 and can populate their physical dimension directly.

Semantic messages can reduce ambiguity during restoration. A fault message should identify the location, confidence, affected protection zone, timestamp and required acknowledgments. An agentic planner can generate candidate switching and resource schedules, but a digital twin must test topology, protection coordination, thermal limits and transient stability. The safety governor should enforce a validity horizon because a plan becomes unsafe if loads, communication links or asset availability change before execution.

A proof-carrying restoration plan would include: the pre-fault and estimated post-fault topology; the assets assumed available; the sequence of switching actions; predicted voltage, frequency and thermal margins; the loads intentionally shed; synchronization conditions; and a certificate that the sequence satisfies the current policy. The plan is rejected if any assumption cannot be verified. The complete sequence, from the fault to the certified execution and the governed update, is traced layer by layer in Section 3.5.3.

### *5.3. Semantic Demand Response and Human Preference Translation*

Demand response traditionally maps price or reliability signals to load changes; the mechanisms are summarized in the overviews of Albadi and El-Saadany [194] and Siano [195], while autonomous game-theoretic consumption scheduling was introduced by Mohsenian-Rad et al. [196] and game-theoretic methods for microgrids, demand-side management and smart-grid communications were reviewed by Saad et al. [197]. Semantic demand response broadens the input from numeric prices to intentions, constraints and social priorities. A household may request “charge the vehicle by 07:00 at minimum cost”, while a clinic may require “never curtail refrigeration and preserve two hours of battery reserve”. An LLM or structured parser can translate such statements into optimization constraints, but the translation must be confirmed and logged.

The main benefit is the preservation of meaning between human objectives and machine schedules. Recent semantic risk-aware VPP work illustrates how unstructured

warnings can pre-activate demand response [167], while the RL-based demand-response literature reviewed by Vazquez-Canteli and Nagy provides the underlying flexibility models and modeling techniques [198]. An agent should distinguish hard requirements, preferences and explanatory text. It should detect conflicts, request clarification and show the operational consequence of each choice. For community mini-grids, this supports transparent prioritization during scarcity and can encode fairness constraints.

Semantic signals can also improve communication efficiency. Event urgency can determine when a device transmits, following the data-significance perspective on semantic communication [166] and the age-of-information measure of freshness [25]. Task-oriented codes can communicate whether flexibility is available without revealing detailed occupancy patterns, in the spirit of task-oriented communication [24] and of the event-driven microgrid semantics of Diao et al. [27]. This creates a connection between demand response, privacy and semantic control.

### *5.4. Predictive Maintenance and Multimodal Asset Intelligence*

Mini-grid maintenance data are naturally multimodal: electrical time series, thermal images, acoustic signals, alarms, work orders, manuals and operator notes. A cognitive maintenance agent can retrieve asset-specific documentation, combine it with anomaly scores, estimate urgency and recommend an inspection. The digital twins examined in Section 4.5 provide the asset history and can evaluate the consequence of delaying maintenance; prognostics and health management are among the industrial twin applications identified by Tao et al. [61], maintenance is among the use cases discussed by Fuller et al. [138], hybrid physics-based and data-driven modeling supports such assessments [141] and the laboratory-scale grid twins of Islam et al. show what real-time synchronization with physical assets requires [142].

Foundation models could eventually provide reusable representations across asset types, but the high-stakes decision should remain decomposed. A perception model identifies evidence; a causal or physics-informed model of the kind proposed in Section 7.4 estimates failure mechanisms; an LLM summarizes and retrieves procedures; and a maintenance planner schedules work subject to spare parts, weather and service continuity. Explanations should expose the measurements and rules that support the recommendation rather than provide only natural-language confidence.

Privacy and data ownership are significant in federated maintenance across multiple mini-grids. FL can share model updates instead of raw data [16] and can accommodate heterogeneous clients through a proximal term [170], while secure aggregation [199] and differential privacy [200] reduce disclosure; federated foundation models extend the same principle to multimodal grid data [125]. However, privacy noise may reduce the detectability of rare faults, creating a trade-off that must be measured rather than assumed.

### *5.5. Community Energy Markets and Inter-Mini-Grid Federations*

Peer-to-peer markets allow prosumers and mini-grids to trade energy, flexibility or reserve, as reviewed for connected communities by Tushar et al. [201] and for peer-to-peer and community-based market designs by Sousa et al. [202]. Game-theoretic incentives were proposed for energy trading among interconnected microgrids [203], distributed optimization was applied to multiple-microgrid energy trading [204] and bilateral contract networks were formulated for peer-to-peer trading [205]. A blockchain-based local energy market was demonstrated in the Brooklyn Microgrid [206], blockchains were proposed for decentralized optimization of microgrid resources [207] and the wider opportunities and challenges of blockchain in the energy sector were systematically reviewed in [208].

Agentic negotiation can express richer offers than price–volume pairs, including carbon origin, critical-load support, reserve duration and willingness to share during emergencies. This flexibility introduces governance questions that are taken up in Section 7.5: which commitments are legally binding, how agents authenticate authority, how collusion is detected and how a community overrides automated trading during scarcity. Semantic interoperability is a prerequisite because the same term, such as "available reserve", must have a consistent physical and contractual interpretation.

Inter-mini-grid federations also create a learning opportunity. Sites can share models, disturbance patterns and cyber indicators without sharing raw customer data. Grid foundation models [52] and federated foundation models [125] may ultimately support cross-site adaptation, but heterogeneity in topology, climate, regulation and hardware requires personalization and uncertainty-aware transfer. Meta-learning-based personalized FL [209] and hierarchical clustering of client updates [210] are established techniques for such personalization.

### *5.6. Cross-Application Synthesis*

The application comparison in Table 26, whose last column lists representative papers for each application class, indicates that no high-value use case is enabled by a single technology. Real-time balancing, as examined in Section 5.1, requires forecasts, distributed coordination and optimization, but cognitive added value appears when flexible resources can express capability, urgency and confidence semantically and when the resulting schedule is independently validated. Resilience and restoration, as examined in Section 5.2, require fault localization, shared topology knowledge, sequence planning and transient checking; they are therefore the application most sensitive to incomplete acknowledgments and stale common knowledge. Semantic demand response, as examined in Section 5.3, depends on translation of human preferences into machine constraints, which makes provenance, privacy and the distinction between hard and soft requirements central. Predictive maintenance, as examined in Section 5.4, benefits from multimodal retrieval and causal diagnosis, although hallucinated procedures or weak document provenance can create new operational hazards. Community markets add negotiation and coalition formation, together with fairness, collusion and contractual ambiguity.

**Table 26.** Application taxonomy for cognitive power mini-grids.

| Application | Cognitive Function | Physical Value | Principal Risk | Representative Papers |
|---|---|---|---|---|
| Balancing and scheduling | Semantic flexibility exchange, distributed planning and tool-based optimization | Lower cost, greater renewable utilization and improved reserve coordination | Oscillatory or unsafe actions from inconsistent beliefs | [119,121–123,161,163,177] |
| Resilience and restoration | Fault semantics, common knowledge and twin-based sequence testing | Faster isolation, restoration and black start | Stale topology, incomplete acknowledgements or invalid transient model | [76,187,188,190–192] |
| Demand response | Human-intent translation and preference-aware optimization | Flexible demand with better participation and fairness | Misclassification of hard constraints or privacy leakage | [166,167,194–196,198] |
| Predictive maintenance | Multimodal retrieval, causal diagnosis and agent scheduling | Reduced failures and maintenance downtime | Hallucinated procedures, weak provenance or domain shift | [16,61,125,141,142,170] |
| Community markets | Semantic offers, negotiation and federated coordination | Local energy sharing and new flexibility services | Manipulation, collusion, contract ambiguity and unfair allocation | [125,201,202,205,206,208] |

The relative maturity of the applications is also different. Balancing and scheduling can build on mature optimization and EMS platforms and are the most immediate candidates for bounded cognitive assistance. Predictive maintenance can initially remain advisory,

which reduces direct physical risk and makes it suitable for early multimodal-agent trials. Demand response and community markets require stronger human and regulatory governance because the decisions affect privacy, equity and contractual commitments. Autonomous restoration offers potentially large resilience gains but should be treated as a later-stage target because an incorrect sequence can violate protection, synchronization or stability requirements. Table 26 therefore supports a staged deployment strategy: begin with advisory diagnosis and certified scheduling; proceed to bounded negotiation; and only then permit autonomous switching after HIL, incident drills and proof-carrying validation.

## 6. Challenges

In this section, an in-depth cross-layer analysis of the barriers that must be resolved before cognitive mini-grids can be deployed safely is provided. Stability, timing contracts and multi-rate interaction; hallucination, grounding failures and tool misuse; semantic interoperability and common-knowledge failure; edge-resource limitations, privacy and sustainability; and cybersecurity, accountability and certification are examined in turn, after which the interactions among these barriers are synthesized. Particular emphasis is placed on how informational or organizational failures can propagate into physical consequences and on the measurable engineering responses that can be introduced across the complete decision lifecycle. A high-performing model is therefore treated as only one component of a safe cognitive system.

### *6.1. Stability, Timing and Multi-Rate Interaction*

Mini-grid control spans microseconds to hours; the corresponding time-scale separation of primary, secondary and tertiary functions is documented in the hierarchical-control literature [5,6]. Protection and converter loops require deterministic latency, while forecasting and planning tolerate slower computation. LLMs and digital-twin searches belong to the slower layers. A cognitive architecture must define timing contracts: maximum observation age, planning deadline, certificate validity and rollback time. The need for such contracts follows from the information-rate limits of feedback control [26,70] and from the tail-risk view of ultra-reliable low-latency wireless communication [211]. Without these contracts, an otherwise correct plan may be executed after its assumptions expire.

Distributed agents can also create interaction instability. Each policy may be stable locally, yet simultaneous set-point changes can produce oscillation. Communication delay, packet loss and asynchronous learning further complicate proofs. Stability analysis must include the semantic encoder and decision scheduler, not only the plant controller. Event-triggered communication must rule out Zeno behavior and preserve sufficient updates during transients.

A practical response is layered runtime assurance: certified local controllers maintain an invariant safe set, learned policies propose reference changes within bounded envelopes, agentic plans are executed in stages and monitors trigger fallback when the state deviates from the validated region. The elements of this response have precedents in the safe-RL literature. Specifically, the safe-RL taxonomy [64], safe model-based RL with stability guarantees [212], constrained policy optimization [213], Lyapunov-based safe RL [164] and shielding [150].

### *6.2. Hallucination, Grounding and Tool Misuse*

LLMs, including those inside the grid-oriented agent systems examined in Section 4.4, can produce fluent but incorrect states, rules, calculations or procedures; hallucination and reasoning limitations are acknowledged in the technical report of a frontier model [58] and in the survey of LLMs by Zhao et al. [48]. Retrieval, as in RAG [30], reduces but does not

eliminate this risk. A retrieved document may be outdated, a tool call may use wrong units or an agent may select a valid solver for the wrong topology, failure modes that the survey of LLM-based agents by Wang et al. [168] and the smart-grid agentic review [35] discuss as open challenges. Every material claim should therefore be grounded in an authoritative source or machine-readable state and every numerical conclusion should be generated by a domain tool rather than language-model arithmetic.

Tool permissions must follow the principle of least privilege. A planning agent may read measurements and run simulations but should not operate breakers. An executor should accept only a signed, schema-valid and certified command. Prompt injection is especially dangerous when external text, maintenance records or market messages can influence the agent. Untrusted content must be separated from system policies and tool schemas.

Evaluation should measure hallucination by consequence. A harmless wording error differs from an invented asset state or missing constraint. Metrics should include the rate of unsupported factual claims, invalid tool call rate, unit-consistency errors, unsafe-plan proposal rate and the fraction caught by validation. Broader research provides taxonomies and diagnostic tools. Hallucination in natural-language generation has been surveyed [214]; concrete AI-safety problems [63] and unsolved machine-learning safety problems [65] have been cataloged; and safe learning, from learning-based control to safe RL, has been reviewed for robotics [215]. Explainability research offers taxonomies [216,217], attribution methods [218,219] and reviews of deep-network explanation [220]. RL for power-system decision and control has likewise been reviewed, with attention to its past and its prospects [221]. Causal responsibility and closed-loop physical consequences must nevertheless be evaluated explicitly in power systems.

### *6.3. Semantic Interoperability and Common-Knowledge Failure*

Current standards provide syntactic and domain semantics for devices, measurements and energy services, but cognitive agents require additional concepts: intent, confidence, urgency, authority, validity horizon, evidence and commitment state. A semantic layer must map these concepts to established models such as the IEC 61850 information model for DERs [156], the CIM of IEC 61970 and IEC 61968 [157], whose practical introduction is given by Uslar et al. [222], the IEEE 2030.5 smart energy profile [223], the OpenADR demand–response profile [224] and the MQTT messaging protocol [225]; the surrounding smart-grid communication technologies and standards are surveyed by Gungor et al. [226] and the communication-infrastructure requirements by Yan et al. [227].

Ontology alignment is difficult across vendors and communities. The same battery may be represented by an inverter controller, battery-management system, market platform and digital twin. Identity, units and temporal scope must be consistent. Semantic compression also risks hiding information needed for diagnosis, which is the principal reservation recorded for the semantic studies examined in Section 4.3; therefore, the receiver must be able to request raw evidence when confidence is low.

Common knowledge can fail because of asymmetric packet loss, malicious acknowledgments or a different interpretation of the same event; the four epistemic states distinguished in Section 4.7.2 make such failures visible. Protocols should expose epistemic uncertainty. Rather than declaring "all agents know the islanding plan", the system should track which agents acknowledged, which evidence they received and whether any critical role remains uncertain.

### *6.4. Resource, Privacy and Sustainability Constraints*

Remote mini-grids may have low-cost controllers, intermittent connectivity and limited technical support. Cloud-dependent agents are unsuitable as the sole decision layer. SLMs, quantized models, sparse event-driven inference and neuromorphic hardware may reduce latency and energy consumption, but capability must be matched to the task; the design space of wireless edge intelligence is surveyed by Park et al. [67], that of edge LLM design and execution is reviewed by Zheng et al. [66] and event-driven neuromorphic semantics were demonstrated in microgrids by Diao et al. [27]. Complex planning can be deferred to a gateway, while device agents use finite-state or symbolic policies.

The AI system itself consumes energy. Green-AI evaluation should report training energy, inference energy, communication energy and operational savings [228]. A model that reduces fuel consumption but requires continuous high-power cloud inference may still be beneficial, but the trade-off must be explicit.

Privacy constraints apply to consumption traces, household preferences, bids, maintenance data and incident logs. Privacy-preserving distributed deep learning [229], FL in mobile edge networks [230], secure aggregation [199] and differential privacy [200] provide partial solutions, but they do not solve model inversion, poisoning, membership inference or regulatory accountability, which remain open problems of federated learning [17].

### *6.5. Cybersecurity, Accountability and Certification*

Cognitive mini-grids enlarge the attack surface beyond the threat classes surveyed in Section 4.7.3. Adversaries can inject false measurements, spoof semantic messages, poison federated updates, manipulate prompts, compromise tools, replay certificates or induce agents to trade against community interests. Smart-grid cybersecurity research already demonstrates the impact of such attacks. False-data injection against state estimation was shown to evade conventional bad-data detection [231] and malicious data attacks were characterized by the regimes in which they remain unobservable [232]. Data-integrity attacks were analyzed in relation to grid topology [233] and the cyber-physical security of the smart grid [182,183] and of cyber-physical systems, in general [184], has been surveyed. Agentic systems add semantic and workflow attacks that these studies did not consider. Foundational work on secure control for survivable cyber-physical systems [185] and on SCADA vulnerability assessment [186] remains directly relevant to the design of these trust boundaries.

Security must be integrated with agent identity and authority. Messages need authentication, integrity, freshness and role-based authorization. The identities of those who proposed, validated, approved and executed each action must be recorded by the system. Standards such as IEC 62351 [151], NISTIR 7628 [152] and NIST SP 800-82 [158] provide a baseline. AI governance frameworks add model-risk and lifecycle controls. Specifically, the NIST AI Risk Management Framework [55], the ISO/IEC 42001 AI management-system standard [160], the ISO/IEC 23894 AI risk-management guidance [56] and the risk-based obligations of the EU AI Act [234]. The cognitive extensions that these assets require are specified in Section 7.6.4.

Certification remains an open problem because learning systems change after deployment. A credible approach certifies the architecture, boundaries, monitors and update procedure rather than attempting to prove every internal neural computation. Updates should be staged through offline testing, digital-twin evaluation, HIL, limited pilot deployment and monitored release.

### *6.6. Cross-Layer Synthesis*

Table 27, whose last column lists representative papers for each challenge, shows that the principal barriers are cross-layer failure chains rather than isolated model defects, consistent with the evidence gap exposed by the cross-study comparison reported in Section 4.7. Timing and stability, as analyzed in Section 6.1, can be degraded by slow reasoning, delayed messages or several locally reasonable agents acting simultaneously. Hallucination, as analyzed in Section 6.2, becomes physically relevant when an incorrect statement is converted into a tool call or set point. Semantic mismatch can create different interpretations of the same reserve or topology state, while common-knowledge failure allows only part of a coalition to execute a coordinated plan; both are analyzed in Section 6.3. Privacy and resource constraints, as analyzed in Section 6.4, can force compression or local processing, but overly aggressive compression may remove evidence needed for safety. Cybersecurity and governance span the entire lifecycle because identity, authorization, model provenance, prompt integrity, action signing and incident accountability must remain connected.

**Table 27.** Major challenges and engineering responses.

| Challenge | Failure Mechanism | Required Response | Key Metric | Representative Papers |
|---|---|---|---|---|
| Timing and stability | Delayed or interacting actions invalidate controller assumptions | Multi-rate architecture, timing contracts and fallback controllers | Worst-case latency, stability margin and rollback time | [6,26,70,150,164,211] |
| Hallucination and tool misuse | Incorrect state, rule, unit or solver invocation | Grounding, schemas, least privilege and independent numerical validation | Unsupported claims and unsafe-plan escape rate | [30,33,35,48,168,214] |
| Semantic mismatch | Agents interpret intent, units or validity differently | Shared ontology, provenance and raw-evidence escalation | Semantic task success and interpretation disagreement | [21,22,156,157,222,223] |
| Common-knowledge failure | Partial reception or false acknowledgement | Epistemic state tracking and quorum/role-aware protocols | Critical-agent acknowledgement and belief divergence | [15,25,114,226,227] |
| Privacy and resources | Data leakage or excessive edge/cloud cost | FL, secure aggregation, SLMs and energy-aware scheduling | Privacy budget and compute/communication energy | [17,66,67,199,200,229,230] |
| Cybersecurity and governance | Spoofing, poisoning, prompt injection or unauthorized actuation | Identity, signed actions, audit, segregation and incident response | Attack success, detection time and attribution completeness | [55,151,152,182,183,231,234] |

The engineering responses in the table should be implemented as mutually reinforcing controls. A multi-rate architecture limits the effect of slow or uncertain intelligence on fast dynamics. Typed schemas and independent solvers reduce hallucination and tool misuse. Ontologies, provenance and raw-evidence escalation reduce semantic ambiguity. Role-aware acknowledgments and abort rules reduce common-knowledge failure. FL, SLMs and energy-aware scheduling reduce exposure and resource cost, while signed actions, network segregation and audit support cyber resilience. The Key Metric column is equally important: each response must be tested against a measurable failure rate rather than described as a design intention. A secure architecture, for example, should report attack success, detection and containment; a grounded agent should report unsupported claims and the unsafe-plan escape rate; and a common-knowledge protocol should report belief divergence and the response to missing critical acknowledgments.

## 7. Future Directions

In this section, the research agenda required to move cognitive power mini-grids from isolated demonstrations toward dependable deployment is developed. Proof-carrying agent actions with runtime assurance; resource-constrained autonomy through edge SLMs;

cross-site knowledge development through federated multimodal foundation models and continual learning; causal digital twins with neuro-symbolic control; and human-agent governance for cognitive energy communities are presented as the principal directions. The agenda is then translated into benchmark scenarios, physical and communication metrics, intelligence and governance measures, reproducibility levels, staged deployment milestones and standards-alignment requirements. Through this progression, emerging technical opportunities are connected to the evidence, safeguards and institutional controls that must be established before greater operational authority is permitted.

### *7.1. Proof-Carrying Agent Actions and Runtime Assurance*

The challenges above converge on one requirement: a high-level action must carry its own evidence of admissibility before it is allowed to reach the plant. Proof-carrying code attaches a machine-checkable proof that a program satisfies a receiver-defined safety policy [73]. The analogous mini-grid concept is a proof-carrying action (PCA). A PCA is defined by Equation (5):

$$\mathcal{P} = \langle a_t, \mathcal{A}, \mathcal{C}, \hat{x}_{t:t+H}, \tau, \varphi, \sigma \rangle. \tag{5}$$

In Equation (5), the proposed action is denoted by $a_t$, its assumptions by $\mathcal{A}$, its checked constraints by $\mathcal{C}$ and its predicted trajectory by $\hat{x}_{t:t+H}$ over prediction horizon $H$. Its validity horizon is denoted by $\tau$, its proof or certificate by $\varphi$ and the identities or signatures of the responsible agents by $\sigma$. The evidence package ($e_t$) of Equation (3) is the part of $\mathcal{P}$ that accompanies the action. Precisely, $\langle \mathcal{A}, \hat{x}_{t:t+H}, \tau, \varphi, \sigma \rangle$. Topology validity, electrical limits, reserve, protection interlocks, solver status, action authority, signatures, state freshness and the validity horizon must be checked by the verifier. The planner does not need to be trusted; only the certificate checker and the underlying models must be trusted, in direct analogy with the small, trusted verifier of proof-carrying code.

Research is required on certificate languages for power systems, rapid verification, uncertain models and sequential actions. Certificates may combine exact checks for topology and protection with probabilistic bounds for forecasts. The safety policy should include electrical limits, reserve requirements, cyber authorization, market commitments and human-approval thresholds.

The certificate $\varphi$ of Equation (5) is specified here as a conjunction of checkable predicates, each with a named checker and its acceptance is specified as five conditions that the verifier evaluates in order. The predicates are the safety properties that the action must preserve: for a switching and dispatch sequence, they are the voltage band on every bus of the affected sections, the loading limit of every branch, the reserve policy after the dispatch, the interlocks of the involved switches, the existence of protection settings for the resulting topology and the authority of the proposing agent for each asset it commits. The assumptions ($\mathcal{A}$) under which the planner established $\varphi$ are the topology version; the load and generation intervals over the horizon; the storage state of charge, along with its tolerance; and the model version of the twin. The verifier accepts the certificate only if the following five conditions hold.

AC1, authority: The signatures ($\sigma$) are valid and the proposing agent holds the authority required for every asset in $a_t$.

AC2, freshness: Each evidence item has an authenticated timestamp and an explicit validity interval on a common clock basis. Its validity, including allowance for clock skew, must cover the entire authorized execution interval and the state-estimate age must remain below the corresponding contract threshold. A continuing dispatch requires full-duration evidence or renewed authorization before each preceding authorization expires.

AC3, assumption consistency: The admissible state and disturbance sets are contained in the assumptions ($\mathcal{A}$) used by the certificate; membership of the point estimate alone is insufficient when estimation error is nonzero.

AC4, independent robust re-evaluation: The verifier recomputes every predicate of $\varphi$ with its own model and with the uncertainty of the estimate and of the model made explicit, as stated in Equation (6).

AC5, recoverability: For every step ($k$) of the sequence, the certificate contains a recovery action ($r_k$) that has itself been certified from the intermediate state reached after step $k$.

The requested plan is accepted only when all applicable conditions hold; otherwise, the failed conditions and their margins are returned to the planner. A reduced plan or a lower authority request is treated as a new candidate whose assumptions, dependencies and recovery actions must be checked again.

The robust re-evaluation of AC4 makes the uncertainty assumptions explicit. Let $\mathcal{X}_t = \{x : (x - \hat{x}_t)^\top \Sigma_t^{-1} (x - \hat{x}_t) \le \kappa^2\}$, where $\Sigma_t$ is positive definite and $\kappa$ is calibrated for a stated coverage probability. A singular covariance requires a separately specified admissible set; covariance alone does not determine coverage. Let $\mathcal{W}$ bound the disturbance ($w_t$) in Equation (1) and let $u_t^a$ denote the candidate low-level input obtained from $a_t$ before authorization. The verifier model ($\hat{f}(x, u, w)$) approximates the state transition ($f$) with a uniform discrepancy bound ($\varepsilon_m$) over the declared operating region. For each physical margin predicate ($\phi_j$), assume a verified Lipschitz constant ($L_j$) in the chosen output norm and a policy margin ($m_j$). The one-step physical check is

$$\inf_{x \in \mathcal{X}_t,\, w \in \mathcal{W}} \phi_j\Big(\hat{f}(x, u_t^a, w)\Big) \;\ge\; L_j\, \varepsilon_m + m_j \qquad \text{for every physical predicate } j, \tag{6}$$

where $L_j \varepsilon_m$ and $m_j$ have the same units as the predicate margin. A deterministic guarantee is conditional on the true state and disturbances remaining in their sets and on the discrepancy and Lipschitz bounds holding. With probabilistic sets, their coverage and dependence must be reported; the result is not an unconditional safety guarantee. The initial admissible state set must also satisfy the current physical safety constraints; otherwise, only a separately verified recovery policy is admissible. For a sequence, this check must be propagated over every reachable intermediate state and relevant intersample interval, with uncertainty accumulated over the prediction horizon. Boolean interlocks, signatures and authorization are checked separately. Increasing uncertainty enlarges the minimization domain and can reduce its worst-case margin; it does not necessarily increase the policy margin ($m_j$). For affine models and affine predicates, extrema over an enclosing interval box occur at vertices; the resulting box minimum is a conservative lower bound for an ellipsoidal set, not generally its exact minimum. For nonlinear models, certified interval or global bounds are required for a robust claim. The minimum of finitely sampled values is generally an upper bound on the true minimum and cannot be substituted as a guaranteed lower bound. Sampling can support only a separately justified statistical test with its confidence and residual risk stated.

The two submissions in Table 28 illustrate the checks on the restoration proposal of step 6 in Table 13. The numerical bounds are stipulated illustrative values, not outputs of a feeder simulation performed in this survey. The first submission fails the freshness condition and two physical predicates. Its requested expiry of 10 min exceeds the 5 min evidence horizon. The 90 min dispatch also requires continuing evidence and renewed authorization. For the scalar voltage check, $L_j = 1$, $\varepsilon_m = 0.01$ per unit and a policy margin of 0.005 per unit require 0.015 per unit above the 0.95 per unit lower bound. The stated voltage of 0.957 per unit leaves 0.007 per unit and therefore fails by 0.008 per

unit. The lower stored-energy bound is $240(1 - 0.05) = 228$ kWh. A dispatch of 150 kW for 90 min consumes 225 kWh, leaving 3 kWh before losses. The 30 min reserve for 180 kW of critical load requires 90 kWh; the additional 10 min policy margin listed in the table adds 30 kWh, giving a total requirement of 120 kWh and a shortfall of at least 117 kWh. The 5 percent loading margin reduces the 90 percent loading limit to 85 percent before any loading-model error allowance.

The revised illustrative submission requests an expiry of 4 min, dispatches B2 at 100 kW for 60 min, sheds 60 kW of deferrable load and connects the mobile unit at bus 14 before closing the tie switch. The stipulated voltage of 0.968 per unit gives a margin of 0.018 per unit against the required 0.015. The lossless energy calculation leaves 128 kWh against the 120 kWh requirement. These two calculations and the stipulated loading of 62 percent satisfy their stated scalar thresholds. They do not establish all certificate conditions: the 60 min dispatch requires continuing valid authorization, uncertainty coverage and model bounds must be justified and recovery and physical interlocks must be verified. Acceptance at A4 is therefore conditional on completing those checks. If the revised voltage were 0.960 per unit, the voltage check would fail by 0.005 per unit. A reduced plan without the tie-switch closure would require a separate full evaluation before acceptance.

If a certificate is missing, expired or inconsistent before execution begins, the sequence is not started. If it fails during execution, after $k$ of $n$ steps have been applied, canceling the remaining commands is not, in general, a safe action because the intermediate state after step $k$ was certified only as a state that the sequence passes through and not as a state in which the system may remain. The recovery rule therefore distinguishes three cases from the measured intermediate state. If the intermediate state is itself certified as safe, the remaining commands are canceled and the system holds; the isolated zone in Table 23 qualifies only after its isolation, local supply and reserve conditions have been verified. If it is not, the verifier evaluates the certified recovery action ($r_k$) of condition AC5, which returns the system to the pre-sequence state when the reverse steps are certified from the measured state, as reopening SW7 is after a failed transfer, or completes the sequence when the remaining steps are certified from the measured state and the reverse steps are not. If neither holds within the expiry, the fallback controller of Layer 2 is invoked with its own certified envelope and the operator is notified. The distinction matters because cancellation leaves the plant wherever the failure found it, whereas recovery returns it to a state that was certified in advance. The state must be estimated again and a new certificate must be obtained before execution resumes. Any partial execution and the taken recovery action must be retained in the audit log.

The certification sequence and rejection path are summarized in Figure 4. The Certify stage corresponds to the safety governor ($\mathcal{G}$) formulated in Section 2.5 through Equation (3), while the certificate fields are specified by Equation (5).

### *7.2. Edge SLMs and Offline-First Autonomy*

Certification presupposes that the deliberative components remain available when external connectivity is lost. Offline-first autonomy assumes that external connectivity may disappear. Local controllers preserve stability; edge agents continue essential scheduling and fault procedures and cloud agents provide nonessential optimization when available. SLMs can support local natural-language interfaces, document retrieval and structured planning but should be distilled for a defined mini-grid vocabulary and tool set rather than used as general chat models.

**Table 28.** Illustrative certificate arithmetic for the restoration proposal of step 6 in Table 13, showing rejection of the first submission and the further checks required for the revised submission.

| Element | Content of the Certificate as Submitted | Evaluation by the Verifier |
|---|---|---|
| *First submission* | | |
| Action ($a_t$) | Close SW7 with SW3 already open by protection, dispatch storage B2 at 150 kW for 90 min, shed 40 kW of deferrable load and route the mobile storage unit to bus 14; execution at $t_{\text{exec}}$, requested expiry of 10 min | Sequence of four steps with a recovery action per step (AC5); the routing of the mobile unit has no plant effect until the unit connects, so it is certified by its own arrival check |
| Assumptions ($\mathcal{A}$) | Topology version 4 with Z3 isolated; load of the re-energized zone in the interval of 160 to 200 kW at 90 percent forecast probability; photovoltaic infeed in the interval of 80 to 140 kW; stored energy of B2 of 240 kWh with a tolerance of 5 percent; twin model version 12 | AC3 requires containment of the admissible state and disturbance sets and a matching topology version; point-estimate membership alone does not establish it |
| Predicates ($\varphi$) | Voltage of every bus in the re-energized zone within 0.95 to 1.05 per unit; loading of every branch not exceeding 90 percent of its rating; reserve of B2 after dispatch sufficient for at least 30 min of critical load; interlock of SW7 satisfied only if SW3 is open; protection settings exist for the topology with SW7 closed; authority of the substation agent over SW7 and B2 | For the scalar voltage check, $L_j = 1$ and $\varepsilon_m = 0.01$ per unit; the policy margins are 0.005 per unit, 5 percent of rating and 10 min of critical load. Separate unit-consistent loading and reserve error bounds are required. The coverage associated with $\kappa = 3$ must be calibrated. |
| Evidence ($e_t$) and signatures ($\sigma$) | State estimate with covariance from Layer 2 timestamped 40 s before submission; acknowledgments of the SW7 and B2 agents; twin evaluation trace; signature of the substation agent | AC1 requires verification of the signed records and authority. AC2 fails: the requested expiry of 10 min exceeds the 5 min evidence horizon; a 90 min dispatch also requires continuing evidence and authorization |
| Robust re-evaluation | Not part of the submission | The stipulated voltage of 0.957 per unit gives a margin of 0.007 against the required 0.015, so the voltage check fails by 0.008 per unit. The lossless reserve is 3 kWh against 120 kWh, including the 10 min margin. Stipulated loading of 78 percent is below 85 percent after the loading margin; a loading-model error allowance remains necessary. SW3 status requires trusted breaker-position and isolation checks; zero current alone is insufficient |
| Decision | Requested authority A4 | Rejected; AC2 and AC4 returned to the planner with the margins by which they failed |
| *Revised submission* | | |
| Action ($a_t$) | Connect the mobile storage unit at bus 14 after its arrival check, close SW7, dispatch B2 at 100 kW for 60 min and shed 60 kW of deferrable load; requested expiry of 4 min | Four proposed steps requiring certified recovery actions; the mobile contribution is stipulated as 80 kW after connection and requires travel, energy, arrival and connection checks |
| Assumptions and evidence | As in the first submission, with the state estimate refreshed 25 s before submission and the acknowledgment of the mobile storage agent added | A 4 min authorization can fit within fresh 5 min evidence after age and clock-skew allowances, but continuing coverage for the 60 min dispatch, signed authority and uncertainty containment must also be established |
| Robust re-evaluation | Not part of the submission | The stipulated voltage of 0.968 per unit gives a margin of 0.018 against the required 0.015; stipulated loading is 62 percent against 85 after the policy margin; lossless reserve is 128 kWh against 120 including the reserve margin. These scalar checks pass, conditional on justified bounds; the remaining predicates require separate verification |
| Recovery actions ($r_k$) | Disconnect the mobile unit; reopen SW7, certified from the closed-SW7 state; stop the B2 dispatch; restore the shed loads in reverse order | AC5 requires separate verification of each recovery from its measured intermediate state, including topology, protection, reserve and connection conditions |
| Decision | Requested authority A4 | Conditionally admissible at A4 only after all acceptance conditions have been independently verified; the illustrative scalar margins alone do not authorize execution |

Rows in italics are group headings.

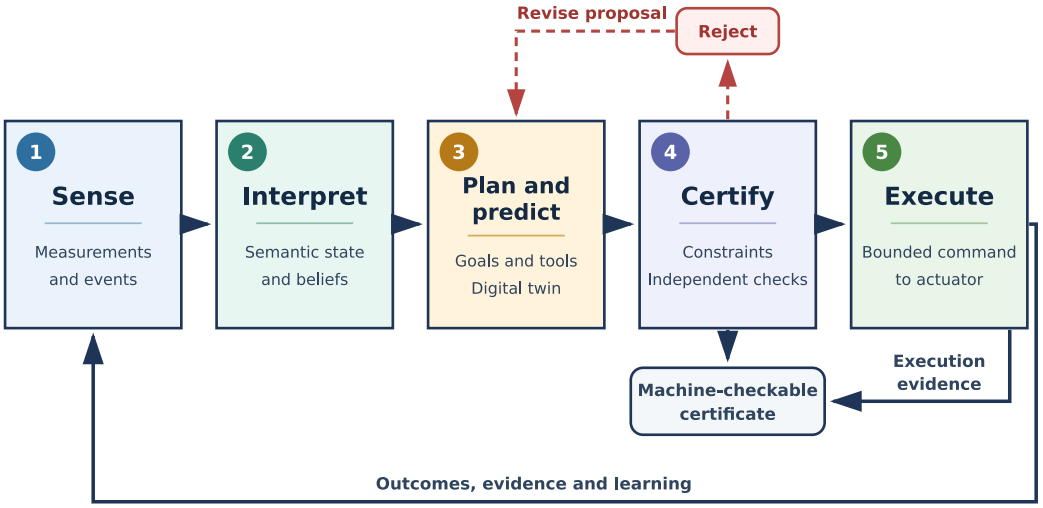


**Figure 4.** Proof-carrying agent action workflow linking sensing, interpretation, planning, independent certification and bounded execution, with rejected proposals revised and execution evidence returned for verification and learning.

SLMs should be compared with symbolic and finite-state alternatives under memory, latency and power constraints in future work. Required metrics include task success per joule, worst-case response time, robustness to corrupted prompts and graceful degradation when retrieval databases are unavailable. Models should be updatable through signed packages and capable of being reverted to a known-good version.

### *7.3. Federated Multimodal Foundation Models and Continual Learning*

Where connectivity exists, the same edge models can benefit from learning across a fleet. A fleet of mini-grids generates heterogeneous data with strong privacy and locality constraints. Federated multimodal foundation models could learn reusable representations for load, weather, topology, asset images and maintenance text. A modular design could use a shared backbone with site-specific adapters [125], physics-aware graph encoders [52] and task heads with calibrated uncertainty [54].

The three emerging sources cited for this direction use nonarchival versions, which qualifies their publication status rather than determining their evidence class. Their reported environments are recorded in Table 21; the proposed integrated direction remains E0 until implemented. The requirements that follow are taken from the peer-reviewed federated-learning literature: the open problems of privacy, robustness and heterogeneity [17]; the systematic challenges of federated learning at scale [124]; and the statistical heterogeneity of client data [170]. Research challenges include non-identically distributed data, missing modalities, communication cost, catastrophic forgetting, poisoning and the difficulty of evaluating transfer to unseen topologies. Continual learning should be triggered by verified distribution shift and should not directly change safety-critical policies. Training data scope, versions, validation conditions and intended tasks must be recorded in a model registry.

### *7.4. Causal Digital Twins and Neuro-Symbolic Control*

Federated models supply hypotheses; the consequences of acting on them must be evaluated before certification. Causal digital twins should answer intervention questions rather than only predict trajectories. Structural causal models [143] can represent how topology, weather, equipment condition and control affect outcomes, while physical equations provide invariant structure; the identification of such models from data is treated by Peters et al. [144] and the learning of causal representations from high-dimensional

observations is reviewed by Scholkopf et al. [145]. Such twins can distinguish correlation from actionable cause and can generate counterfactual explanations for operators.

Neuro-symbolic control complements causal twins by combining learned perception with explicit rules and ontologies [47,169]. Symbolic constraints can enforce voltage, current, reserve, authorization and procedural rules, while neural components adapt to uncertainty, as the grid-oriented neuro-symbolic preprint illustrates [62]. Research should move from post hoc explanations to executable symbolic safety shields and traceable rule activation, consistent with the argument that high-stakes decisions require inherently interpretable models rather than explanations of black boxes [235].

Neuromorphic semantic control is a related direction for very low-power and low-latency devices. SNNs with event-driven data selection can infer coordination information from sparse electrical events, reducing the dependence on a dedicated communication layer [27]. The challenge is to provide stability guarantees and interoperable semantics for learned spike representations.

### *7.5. Human-Agent Governance and Cognitive Energy Communities*

Technical assurance alone is insufficient because autonomy also changes organizational responsibility. Communities and operators need to define which objectives agents can optimize, which loads receive priority, who approves exceptional actions and how decisions are appealed. Natural-language interfaces can improve accessibility but must show constraints, alternatives, uncertainty and expected impact.

Research should study calibrated trust rather than maximum trust. Operators should understand when to rely on the agent and when to intervene. Interfaces can provide contrastive explanations (for example, why one restoration sequence was rejected) and can expose the certificate supporting an accepted action. Human feedback can improve planning, but it must be separated from immutable safety rules.

At ecosystem scale, mini-grids may form the federations examined in Section 5.5 that trade, share reserves and learn collectively. Governance must prevent larger participants from dominating semantic standards or model updates. Fairness, energy justice and local ownership should be treated as design objectives, especially in remote electrification contexts. Deployment in such contexts is shaped by the market conditions documented by the World Bank [41] and by Sustainable Energy for All [236], as well as by the access gap analyzed in the Africa Energy Outlook [237]; the governance of peer-to-peer and community markets raises the same questions [201,202].

The future directions in Table 29 are ordered by the assurance gap they close. Proof-carrying actions address the immediate trust boundary between a flexible planner and an actuator. Offline-first edge agents, as described in Section 7.2, address continuity when cloud services, wide-area communication or external model endpoints are unavailable. Federated foundation models, as described in Section 7.3, address representation reuse across heterogeneous sites without centralizing all operational data. Causal and neuro-symbolic twins address counterfactual reasoning and explicit constraint enforcement. Neuromorphic semantic control addresses communication and computation overhead at fast edge timescales. Human-agent governance addresses authority, accountability and social acceptance across all of the preceding technologies.

The Minimum Experimental Evidence and Deployment Milestone columns prevent these directions from remaining visionary labels. A proof-carrying method should report certificate construction and verification time, as well as rejection behavior and independent verifier coverage under faulted cases. An edge SLM should be tested under link loss, stale local knowledge and constrained hardware rather than only on a desktop benchmark. A federated foundation model should include multi-site evaluation with non-identically dis-

tributed data, privacy tests, poisoning tests and model-registry governance. A causal twin should be evaluated through interventions and out-of-distribution events, not only forecast accuracy. Neuromorphic semantics should report closed-loop stability and hardware energy, in addition to reduced packet count. Human-agent systems should use representative participants, incident exercises and auditable decision records. These requirements convert the future agenda into a sequence of falsifiable research programs.

**Table 29.** Future research directions and minimum evidence needed for progress.

| Direction | Research Objective | Minimum Experimental Evidence | Deployment Milestone |
| --- | --- | --- | --- |
| Proof-carrying actions | Verify high-level plans without trusting the planner | Faulted and uncertain cases with certificate timing and rejection analysis | Independent runtime verifier on HIL platform |
| Offline-first edge agents | Maintain useful cognition without cloud access | Power, memory and latency tests under link loss and stale databases | Signed, rollback-capable edge deployment |
| Federated foundation models | Transfer multimodal representations across heterogeneous sites | Multi-site evaluation with non-identically distributed data, privacy tests and poisoning tests | Governed model registry and site-specific adapters |
| Causal/neuro-symbolic twins | Produce counterfactuals and enforce explicit constraints | Intervention tests, rule traces and out-of-distribution scenarios | Closed-loop twin with bounded authority |
| Neuromorphic semantics | Reduce communication and computation with event-driven meaning | Stability, bandwidth and hardware-energy measurements | Interoperable semantic event profile |
| Human-agent governance | Align autonomy with community and operator authority | Controlled user studies, incident drills and audit completeness | Approved operating procedures and accountability model |

### *7.6. Evaluation Benchmarks, Metrics and Deployment Roadmap*

The future directions are completed by an evaluation framework that covers benchmark scenarios and platforms, metrics, reproducibility and evidence levels and the roadmap with its standards alignment. Evidence across physical, communication, intelligence, assurance and human dimensions is required instead of isolated accuracy or cost comparisons. As a survey, the contribution provided here is a proposed evaluation specification with implementation requirements. Empirical validation of the proposed architecture through simulation, experiments or dataset analysis is reserved for future benchmark studies.

#### 7.6.1. Benchmark Scenarios and Platforms

A cognitive mini-grid benchmark that closes the evidence gap synthesized in Section 4.7.5 should include at least five scenario families: normal grid-connected scheduling, islanded renewable variability, communication degradation, compound cyber-physical faults and community scarcity. Each family should specify topology, asset parameters, measurement noise, communication delays, agent roles, policies and expected safety invariants. The communication-degradation and compound-fault families are intended to reproduce the cross-layer failure chains synthesized in Section 6.6.

Simulation platforms such as MATPOWER [238], GridLAB-D [239], OpenDSS [240] and PyPSA [241] can support steady-state and scheduling studies, while electromagnetic transient and real-time digital simulators are required for converter and protection behavior. The power simulator should be connected to communication and agent environments

through co-simulation. HIL and power-HIL should be required before claims of real-time deployment.

Existing openly shared resources should be used for reproducible evaluation when their scope is suitable. SimBench provides benchmark grid models and year-long load, generation and storage profiles [242]; IEEE Power and Energy Society test feeders provide standard distribution topologies [243]; and the public End-Use Load Profiles dataset provides 15 min building-load series across building types and climate regions [244]. These resources do not validate the architecture proposed in this survey. They are identified as suitable inputs for future confirmation studies, in which the selected topology, profile subset, preprocessing, license, version and access date must be reported.

Table 30 is a reporting template rather than a record of experiments performed in this survey. Its use would prevent a future evaluation from claiming reproducibility without identifying the public data, physical assumptions, communication conditions, algorithms and statistical procedure.

**Table 30.** Minimum simulation and dataset parameters to be reported in a cognitive mini-grid evaluation.

| Parameter Group | Required Entries | Units or Reporting Form |
|---|---|---|
| Dataset and topology | Resource name, version, license, feeder identifier, buses, branches and phase model | Identifier, count and access date |
| Time-series data | Load, renewable generation, prices, weather, sampling interval, missing-data treatment and split rule | kW, kWh, currency/kWh, seconds or minutes and percentages |
| Energy assets | Rated power and energy, efficiency, state-of-charge bounds, ramp limits and degradation model | kW, kWh, percentage, kW/s and stated model |
| Electrical limits | Nominal voltage and frequency, voltage band, line limits, reserve and protection assumptions | V or per unit, Hz, A or kVA and percentage |
| Communication | Delay, jitter, loss, bandwidth, update rule, acknowledgement timeout and attack model | ms, percentage, bit/s, events/s and scenario definition |
| Learning and agents | Observations, actions, reward or objective, model version, seed, training steps and update policy | Explicit schema, count, version and seed list |
| Optimization and control | Horizon, time step, solver, tolerance, stopping rule, infeasibility policy and fallback controller | Seconds or hours, solver version and tolerance |
| Evaluation | Baselines, repetitions, confidence interval, disturbance cases, ablations and evidence level | Count, percentage and Levels 0 to 5 |

A proposed CognitiveMiniGridGym should expose a reproducible interface rather than only a benchmark name. Its observation space should include electrical state, communication state, ontology version, local beliefs and acknowledged commitments. Its action space should distinguish advice, tool calls and candidate set points. Time steps should be aligned with the control hierarchy, while scenario generators should represent electrical disturbances, communication failures, cyber events and resource scarcity. Reset behavior, deterministic random seeds, expected invariants, complete logs, machine-readable configuration files and public-data provenance should be specified. Baselines should include centralized optimization, distributed optimization, MPC, RL and MARL; semantic communication without LLM agents; and agentic planning without semantic compression. This separation is necessary to determine which component produces each benefit.

The reporting template of Table 30 is instantiated as the proposed reference scenario, RS1, in Table 31. The feeder, asset placements, disturbance sequence and candidate baselines are specified as a starting point for matched evaluation. RS1 is not an experiment performed in this survey. Its remaining implementation choices must be fixed in a machine-readable

configuration before independent results can be compared. The feeder is the 33-bus radial distribution system of Baran and Wu [245], with 32 initially closed branches, 5 normally open tie lines, a nominal voltage of 12.66 kV and nominal total load of 3715 kW and 2300 kvar. It is treated as a community mini-grid with a single point of common coupling at bus 1. Public End-Use Load Profiles [244] and SimBench generation profiles [242] are proposed as inputs to an OpenDSS implementation [240]. The former are calibrated outputs of physics-based models of the United States building stock. Their combination with SimBench photovoltaic profiles forms a constructed scenario whose geographic, weather and temporal compatibility must be declared. No selected profile files or simulation outputs are supplied by this survey.

**Table 31.** Proposed reference scenario RS1 for fault isolation and restoration under communication degradation, with author-selected asset placements and event times.

| Element | Specification |
| --- | --- |
| Feeder | 33-bus radial distribution system [245]: 33 buses; 32 branches; 5 tie lines connecting the bus pairs (8, 21), (9, 15), (12, 22), (18, 33) and (25, 29); nominal voltage of 12.66 kV; total load of 3715 kW and 2300 kvar; sectionalizing switches on every branch; point of common coupling at bus 1 with a 5 MVA connection; author-selected voltage band of 0.95 to 1.05 per unit; branch loading limit of 90 percent of separately specified ratings because the public base case does not supply finite thermal ratings |
| Assets | Photovoltaic units of 500 kW at buses 18 and 25 and 400 kW at bus 33; grid-forming battery of 1 MW and 2 MWh at bus 6 with 70 percent state of charge at the start; battery of 500 kW and 1 MWh at bus 30 at 60 percent; mobile storage unit of 250 kW and 500 kWh parked at bus 1 with a travel time of 20 min to any bus of the feeder; critical loads of 200 kW in total at buses 13, 14 and 24; deferrable load of 15 percent of the load at every bus; reserve policy of 30 min of critical load withheld in the grid-forming battery; every storage unit follows a linear state-of-charge model with a round-trip efficiency of 90 percent, symmetric power limits and state-of-charge limits of 10 and 95 percent |
| Time series | 15-min load profiles of the End-Use Load Profiles dataset [244]: 200 single-family residential and 20 small-office building profiles from the ResStock and ComStock AMY2018 release 1 for the county in ASHRAE climate zone 3A with the largest residential sample, aggregated per bus in proportion to the nominal bus loads of the feeder so that the aggregate peak equals 3715 kW, with interval energies converted to mean power over the 15 min interval before scaling; SimBench photovoltaic profile PV1 [242] for a clear summer day, scaled to the unit ratings; proposed 24 h simulation in local standard time at a 1 s quasi-static physical step and a 1 min agent step, with 15 min profiles held constant within each interval; exact release paths, building IDs, county, dates, PV column and time-zone alignment must be recorded in the input manifest |
| Disturbance sequence | 12:00:00 permanent three-phase fault on the branch between buses 7 and 8, cleared by protection within 100 ms, which is imposed as a scenario event on the quasi-static model rather than simulated electromagnetically, isolating buses 8 to 18; 12:10:00 photovoltaic infeed drops by 40 percent for 30 min; 12:15:00 false-data injection of plus 0.03 per unit on the reported voltage of bus 14 for 5 min; 12:30:00 restoration through healthy ties may be considered after isolation and protection checks; the permanently faulted branch remains open unless a separate repair event is specified |
| Communication model | Message delay of 50 ms plus an exponentially distributed jitter with a mean of 150 ms, giving a mean delay of 200 ms, message loss of 5 percent independent per message, partition of the agent of the switch on branch 14 to 15 from 12:03:00 to 12:04:30, ontology mismatch injected at 12:06:00 by presenting version 3 of the reserve definition to the agent of the battery at bus 30 while every other agent uses version 4, acknowledgment timeout 2 s and commit deadline of 30 s before execution |
| Agent roles | One agent per battery, per photovoltaic unit, per switch and for the mobile storage unit; a substation agent as committer; a planning agent with the digital twin at Layer 5; and a governor at Layer 6 with the acceptance conditions of Section 7.1 instantiated for the RS1 asset identifiers and limits |
| Baselines | BL0, manual restoration with a fixed 30 min crew response and no automation; BL1, rule-based restoration following the dispatch and transition functions of IEEE Std 2030.7 [153] with perfect communication, implemented as a fixed priority list that closes the tie switches in the order 8 to 21, 12 to 22 and 18 to 33, subject to the voltage and loading limits; BL2, centralized restoration by a mixed-integer linear program with perfect communication and full observability that maximizes the critical load restored over the horizon, subject to linearized DistFlow power-flow constraints, storage energy balance and switch interlocks, with a declared solver, tolerance and stopping rule at each 1 min step and the achieved optimality gap reported; BL3, the proposed architecture with the governor removed so that the planning agent actuates directly; BL4, the proposed architecture as specified |
| Metrics | Energy not served in kWh; critical load served in percent of the critical-load energy over the day; time to restore the last critical load in minutes; voltage violations as bus-minutes outside the band; branch overloads as branch-minutes above the limit; unsafe actions blocked by the governor, counted; coordinated actions aborted for lack of common knowledge, counted; message volume in kilobytes and mean age of information in seconds at the substation agent; decision latency of the planning agent in seconds; energy consumed by agent computation in kWh |

**Table 31.** *Cont.*

| Element | Specification |
| --- | --- |
| Repetitions and reporting | Ten runs with seeds 1 to 10 controlling the load noise, which is Gaussian with a standard deviation of 5 percent of each interval load, the message loss and the jitter; means and 95 percent confidence intervals from the t-distribution with nine degrees of freedom reported per metric; every item of Table 30 disclosed; ablations BL3 against BL4 and BL2 against BL4 mandatory |
| Evidence class | E0 for this proposed specification; E2 is the target for a completed closed-loop simulation. E3 additionally requires a documented real-time hardware-in-the-loop experiment with timing, interface and synchronization evidence |

The baselines are intended to separate the effects of automation, planning, communication and the governor. BL0 and BL1 characterize manual and rule-based restoration. BL2 supplies a centralized optimization comparator under its stated linearized model and information assumptions; it is not a demonstrated optimum for the nonlinear physical system. Matched degraded-communication variants of BL1 and BL2 are also required before robustness is attributed to BL4. BL3 against BL4 can test the governor only when planner inputs, data, resources and random realizations are matched. Whether false-data injection causes an unsafe proposal, whether such a proposal is blocked and whether a violation occurs must be measured. The partition and ontology-mismatch events test the preparation and local execution rules of Section 4.7.2; zero inconsistent execution is a validation target, not an established outcome.

Before RS1 can be reproduced, its exact data files and licenses, dates, electrical and storage equations, branch ratings, active and reactive load treatment, mobile-storage initial energy, charging and discharging efficiencies, noise timing and clipping, solver version and tolerances, planner implementation and metric integration rules must be fixed. The stated 90 percent round-trip efficiency must be allocated consistently between charging and discharging. The input manifest must identify the selected building IDs and photovoltaic column. Resampling 15 min profiles to a 1 s grid does not provide measured variability at the finer resolution. The imposed 100 ms protection event cannot validate fault-clearing transients in the quasi-static model. A completed benchmark must report all choices in Table 30 and match tasks, data, physical constraints, information and computation budgets before outcomes are compared.

#### 7.6.2. Metrics

Physical metrics include energy cost, renewable curtailment, unserved energy, voltage and frequency deviation, stability margin, restoration time, battery degradation and emissions. Forecasts should be assessed with explicitly defined errors such as the root mean square error (RMSE) and mean absolute error, while scale-dependent comparisons across heterogeneous series should be interpreted carefully [246]. For $n$ observations ($y_t$) and forecasts ($\hat{y}_t$), RMSE is defined as

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}. \tag{7}$$

Equation (7) penalizes large errors more strongly because the residuals are squared before averaging; it is expressed in the unit of the observations, such as kW for a power forecast. Mean absolute error (MAE) is the arithmetic mean of the absolute residuals in the same unit and is less sensitive to single large errors [246]. Communication metrics include bytes, packets, age of information, deadline violations and energy per useful message. Age of information measures the elapsed time since the generation of the freshest received update rather than only transmission delay [25]. Semantic metrics include task success, intent preservation, disagreement and information sufficiency. Agentic metrics include

planning success, tool-call correctness, unsupported claims, proposed and intercepted constraint violations, explanation fidelity and recovery from failed tools.

Assurance metrics should distinguish proposal safety from execution safety. A system may propose many unsafe actions but reject all of them; this still imposes workload and may indicate poor planning. The unsafe proposal rate, verifier false acceptance, verifier false rejection, certificate time, rollback success and coverage of monitored invariants should be reported. Cyber metrics include attack success, time to detection, containment and attribution completeness. Human metrics include workload, trust calibration and decision time, for which appropriate reliance rather than maximal trust should be evaluated [247]. Privacy budgets should be interpreted according to the employed formal differential-privacy mechanism [200].

Table 32 establishes that a cognitive mini-grid cannot be evaluated by a single aggregate objective. Physical metrics remain primary because lower cost or reduced communication has no value if voltage, frequency, protection or energy service deteriorates. Communication metrics quantify the resource burden and timeliness of coordination. Semantic metrics determine whether the compressed or abstracted message preserved the intent and evidence needed for the task. Learning and agent metrics isolate planning quality, tool correctness and latency. Assurance metrics reveal whether unsafe proposals are being generated, whether the verifier catches them and whether rollback succeeds. Cyber, human and Green-AI dimensions measure attack resilience, accountable use and the net resource cost of the intelligence itself.

**Table 32.** Recommended metric families for cognitive mini-grid evaluation.

| Dimension | Core Metrics | Reporting Requirement |
|---|---|---|
| Physical | Cost, curtailment, unserved energy, voltage/frequency error, degradation and emissions | Mean, tail and worst-case values across disturbances and seeds |
| Communication | Bandwidth, packet rate, age of information, deadline misses and communication energy | Compare against raw-data and event-triggered baselines |
| Semantic | Task success, intent preservation, disagreement and uncertainty calibration | Evaluate under noise, packet loss and ontology mismatch |
| Learning and agents | Forecast RMSE, reward, sample efficiency, tool correctness, hallucination and planning latency | Separate model reasoning time from tool execution and validation |
| Assurance | Unsafe proposals, false acceptance/rejection, certificate time and rollback success | Publish invariant definitions and failure cases |
| Cyber/privacy | Attack success, detection/containment, privacy budget and poisoning robustness | Include adaptive and compound attacks |
| Human and governance | Workload, trust calibration, override rate and audit completeness | Use representative operators or community participants |
| Green AI | Training/inference energy, communication energy and net operational saving | Report hardware, model size and duty cycle |

The reporting requirements are designed to expose trade-offs and tail behavior. Mean cost should be accompanied by worst-case unserved energy and disturbance-specific stability. A semantic method should be compared with both raw-data and conventional event-triggered baselines under the same channel. Agentic latency should be decomposed into model reasoning, retrieval, tool execution and validation so that the bottleneck is visible. An assurance system should report false acceptance and false rejection because an excessively conservative verifier can make the system unusable, even when it is safe. Green-AI reporting should combine training, inference and communication energy with the operational energy saved by the controller. The resulting evaluation is necessarily multi-

objective; Pareto fronts or constrained comparisons are more informative than declaring one approach best from a single weighted score.

#### 7.6.3. Reproducibility and Evidence Levels

Evidence should be classified according to the six levels (E0 to E5) defined in Table 6 and the horizon, authority, scale and communication classes of the same table should be reported alongside it so that a reader can locate the study in the comparison space before any outcome is read. A paper should not generalize simulation results to deployment readiness without explicitly addressing the gap. The same evidence criteria must be applied uniformly to every cited study.

Reproducibility requires source code, configuration files, random seeds, scenario definitions, model versions, prompt templates, ontology versions and tool schemas. LLM studies should record decoding settings and the model endpoint/version. Federated studies should report client heterogeneity, participation, communication rounds, privacy mechanisms and attack assumptions. Safety studies should publish the invariants and counterexamples.

#### 7.6.4. Roadmap and Standards Alignment

The roadmap in Figure 5 uses the single time marker of "2026 and after" for a five-stage research and deployment agenda. Stage 1 establishes foundations through ontologies, semantic metrics and common-knowledge protocols. Stage 2 develops safe prototypes based on digital twins, edge SLMs and proof-carrying actions. Stage 3 establishes interoperable pilots through multi-vendor HIL evaluation, federated sites and human-agent procedures. Stage 4 advances assured autonomy through certification, runtime assurance and autonomous restoration. Stage 5 develops cognitive ecosystems through inter-mini-grid markets, fleet learning and regulatory adoption. The stages indicate an ordered agenda and do not assign separate calendar years to activities after 2026. The roadmap stages are distinct from the capability-maturity levels of Figure 3; their relationship and the evidence and authority admissible at each stage are stated in Table 10. The placement of autonomous restoration in Stage 4 follows the deployment order derived from the application comparison reported in Section 5.6. The normative assets that each stage must satisfy are those mapped to the layers in Section 3.5.2 so that the roadmap advances through the existing standards rather than beside them.

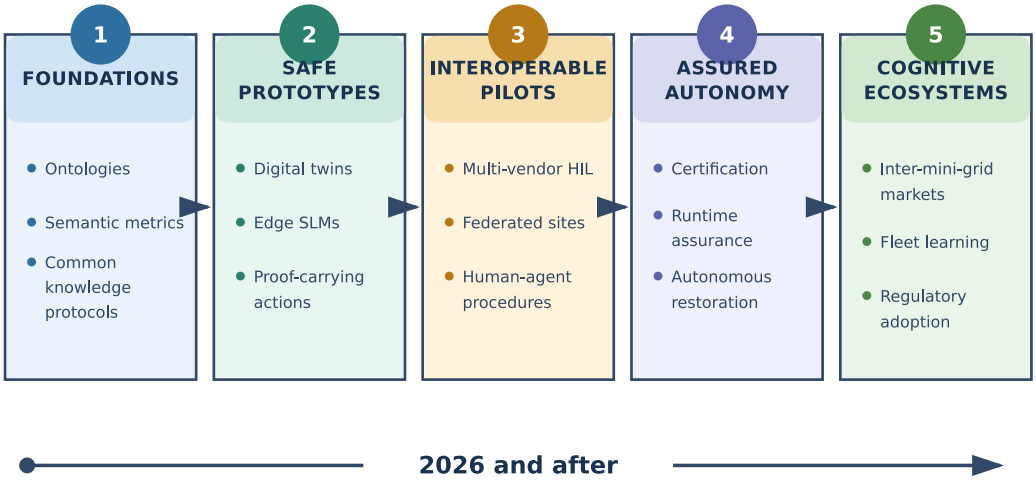


**Figure 5.** Indicative five-stage research and deployment roadmap for 2026 and after, with ordered stages rather than fixed adoption dates.

Standards alignment should begin immediately. Controller functions and testing are defined by IEEE 2030.7 [153] and IEEE 2030.8 [154]. Recommended practice for microgrid planning and design is provided by IEEE 2030.9 [159]. DER interconnection is defined by IEEE 1547 [155]. Energy services are supported by IEEE 2030.5 [223] and OpenADR [224]. Information models are provided by IEC 61850-7-420 [156] and the CIM [157]. Security requirements are provided by the IEC 62351 series [151], NISTIR 7628 [152] and NIST SP 800-82 [158]. The research gap identified in this survey is an agentic semantic profile that adds intent, evidence, confidence, authorization and certificates without creating a competing information model.

Table 33 indicates that the cognitive mini-grid should extend existing interoperability and governance assets rather than create an isolated replacement stack; the complementary question of which layer each asset already governs is answered in Table 12. IEEE 2030.7 and 2030.8 already define microgrid-controller functions and testing, so an agentic profile should add tests for role behavior, semantic failure, certificate verification and fallback. IEEE 1547 already constrains DER interconnection and permitted functions, which means every agent command should map to a defined and authorized DER capability. IEC 61850 and CIM provide the strongest basis for common device and system meaning; the required additions concern intent, confidence, validity, evidence, acknowledgment and commitment. IEEE 2030.5 and OpenADR can carry energy-service and demand-response interactions, but machine negotiation requires explicit authority and audit metadata.

Security and AI-governance assets cover complementary parts of the problem. IEC 62351 and NIST operational-technology guidance address authentication, integrity, segmentation and incident response, while AI risk and management frameworks address lifecycle governance, monitoring and accountability. The missing connection is a power-specific safety case for learning and agentic components. Such a profile should define model and ontology versioning, signed prompts or tool manifests where appropriate, least-privilege tool access, independent action verification, update approval and evidence retention. The regulatory row further indicates that autonomy must be classified by consequence and authority, not by whether an LLM is present. An advisory agent, a scheduler that produces certified set points and an autonomous restoration agent should have different conformity and accountability requirements, even when they share the same underlying model.

**Table 33.** Standards and governance assets relevant to a cognitive mini-grid profile.

| Asset | Current Contribution | Required Cognitive Extension |
|---|---|---|
| IEEE 2030.7 and IEEE 2030.8 | Controller functions and testing | Agent role tests, semantic failure cases and certificate verification |
| IEEE 1547 | DER interconnection and interoperability | Mapping of agent commands to permitted DER functions |
| IEC 61850 and CIM | Device and system information models | Intent, confidence, validity horizon, evidence and commitment semantics |
| IEEE 2030.5 and OpenADR | Energy-service and demand–response exchanges | Machine-negotiated preferences with authority and audit metadata |
| IEC 62351, NISTIR 7628 and NIST SP 800-82 | Authentication, integrity and operational-technology security guidance | Prompt/tool isolation, model provenance and signed agent actions |
| NIST Artificial Intelligence Risk Management Framework, ISO 42001 and ISO 23894 | AI risk and management processes | Power-specific safety cases, monitoring and update governance |
| EU AI Act and local regulation | Risk-based legal obligations | Classification and accountability for autonomous energy decisions |

## 8. Conclusions

The cognitive power mini-grid is best understood as a composition of established control engineering with new semantic, distributed and agentic capabilities; the eight capabilities that define cognition in Section 1 are realized by this composition rather than by any single algorithm. Its novelty is attributed to the coordination of heterogeneous agents through which task-relevant meaning is exchanged, sufficient shared knowledge is established, verified tools are used for reasoning, consequences are predicted and only actions that satisfy physical and governance constraints are executed.

In answer to the four research questions formulated in Section 2.2, five principles were identified. First, deterministic local control and protection remain the foundation. Second, semantics must preserve the evidence required for safety. Third, common knowledge and authority are explicit prerequisites for coordinated actions. Fourth, agentic planning must be separated from execution by independent validation and runtime assurance. Fifth, evaluation must jointly measure physical outcomes, communication efficiency, epistemic consistency, cyber resilience, human governance and AI energy cost.

The most important research directions are proof-carrying agent actions, offline-first edge SLMs, federated multimodal foundation models, causal and neuro-symbolic digital twins, neuromorphic semantic control and accountable human-agent energy communities. Progress should be judged by staged evidence from simulation to HIL and field pilots rather than by conversational demonstrations or single-scenario cost reductions. Under the evaluation framework specified in Section 7.6, that evidence is graded on the unified scale of Table 6, compared only under the comparability rules of Section 2.4, read against the applicability boundaries appraised in Section 4.7.4, measured with the metric families defined in Section 7.6.2 on the benchmark scenarios specified in Section 7.6.1 and sequenced by the five-stage roadmap. With these boundaries, cognitive autonomy can become a disciplined extension of mini-grid engineering rather than an uncontrolled substitution for it.

**Author Contributions:** Conceptualization, I.I. and S.J.; methodology, I.I. and S.J.; investigation and literature synthesis, I.I. and S.J.; taxonomy and architecture development, I.I. and S.J.; writing—original draft preparation, I.I.; writing—review and editing, I.I. and S.J.; visualization, I.I. All authors have read and agreed to the published version of the manuscript.

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

**Data Availability Statement:** No new experimental data were created in this survey. The reviewed literature, standards and public software resources are identified in the references and evaluation roadmap.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

| | |
|---|---|
| 5G | Fifth Generation |
| 6G | Sixth Generation |
| AC | Alternating Current |
| ADMM | Alternating Direction Method of Multipliers |
| AI | Artificial Intelligence |
| AUC-ROC | Area Under the Receiver Operating Characteristic Curve |
| BDIx | Extended Belief–Desire–Intention |
| CIM | Common Information Model |
| COMA | Counterfactual Multi-Agent Policy Gradients |
| DAI | Distributed Artificial Intelligence |
| D2D | Device to Device |

| | |
|---|---|
| DC | Direct Current |
| DER | Distributed Energy Resource |
| DRL | Deep Reinforcement Learning |
| EMS | Energy Management System |
| FL | Federated Learning |
| HIL | Hardware in the Loop |
| LLM | Large Language Model |
| LSTM | Long Short-Term Memory |
| MACKRL | Multi-Agent Common Knowledge Reinforcement Learning |
| MADDPG | Multi-Agent Deep Deterministic Policy Gradient |
| MAE | Mean Absolute Error |
| MARL | Multi-Agent Reinforcement Learning |
| MAS | Multi-Agent System |
| MPC | Model Predictive Control |
| NISTIR | National Institute of Standards and Technology Interagency Report |
| PCA | Proof-Carrying Action |
| PRISMA | Preferred Reporting Items for Systematic Reviews and Meta-Analyses |
| RAG | Retrieval-Augmented Generation |
| RL | Reinforcement Learning |
| RMSE | Root Mean Square Error |
| SCADA | Supervisory Control and Data Acquisition |
| SLM | Small Language Model |
| SNN | Spiking Neural Network |
| SP | Special Publication |
| VDN | Value-Decomposition Network |
| VPP | Virtual Power Plant |

## References


1. Lasseter, R.H. MicroGrids. In Proceedings of the IEEE Power Engineering Society Winter Meeting, New York, NY, USA, 27–31 January 2002; Volume 1, pp. 305–308. [CrossRef]
2. Hatziargyriou, N.; Asano, H.; Iravani, R.; Marnay, C. Microgrids. *IEEE Power Energy Mag.* **2007**, *5*, 78–94. [CrossRef]
3. Katiraei, F.; Iravani, R.; Hatziargyriou, N.; Dimeas, A. Microgrids Management. *IEEE Power Energy Mag.* **2008**, *6*, 54–65. [CrossRef]
4. Guerrero, J.M.; Vasquez, J.C.; Matas, J.; de Vicuna, L.G.; Castilla, M. Hierarchical Control of Droop-Controlled AC and DC Microgrids: A General Approach Toward Standardization. *IEEE Trans. Ind. Electron.* **2011**, *58*, 158–172. [CrossRef]
5. Bidram, A.; Davoudi, A. Hierarchical Structure of Microgrids Control System. *IEEE Trans. Smart Grid* **2012**, *3*, 1963–1976. [CrossRef]
6. Olivares, D.E.; Mehrizi-Sani, A.; Etemadi, A.H.; Canizares, C.A.; Iravani, R.; Kazerani, M.; Hajimiragha, A.H.; Gomis-Bellmunt, O.; Saeedifard, M.; Palma-Behnke, R.; et al. Trends in Microgrid Control. *IEEE Trans. Smart Grid* **2014**, *5*, 1905–1919. [CrossRef]
7. Dragicevic, T.; Lu, X.; Vasquez, J.C.; Guerrero, J.M. DC Microgrids—Part I: A Review of Control Strategies and Stabilization Techniques. *IEEE Trans. Power Electron.* **2016**, *31*, 4876–4891. [CrossRef]
8. Dragicevic, T.; Lu, X.; Vasquez, J.C.; Guerrero, J.M. DC Microgrids—Part II: A Review of Power Architectures, Applications and Standardization Issues. *IEEE Trans. Power Electron.* **2016**, *31*, 3528–3549. [CrossRef]
9. McArthur, S.D.J.; Davidson, E.M.; Catterson, V.M.; Dimeas, A.L.; Hatziargyriou, N.D.; Ponci, F.; Funabashi, T. Multi-Agent Systems for Power Engineering Applications—Part I: Concepts, Approaches and Technical Challenges. *IEEE Trans. Power Syst.* **2007**, *22*, 1743–1752. [CrossRef]
10. McArthur, S.D.J.; Davidson, E.M.; Catterson, V.M.; Dimeas, A.L.; Hatziargyriou, N.D.; Ponci, F.; Funabashi, T. Multi-Agent Systems for Power Engineering Applications—Part II: Technologies, Standards and Tools for Building Multi-Agent Systems. *IEEE Trans. Power Syst.* **2007**, *22*, 1753–1759. [CrossRef]
11. Dimeas, A.L.; Hatziargyriou, N.D. Operation of a Multiagent System for Microgrid Control. *IEEE Trans. Power Syst.* **2005**, *20*, 1447–1455. [CrossRef]
12. Nunna, H.S.V.S.K.; Doolla, S. Multiagent-Based Distributed-Energy-Resource Management for Intelligent Microgrids. *IEEE Trans. Ind. Electron.* **2013**, *60*, 1678–1687. [CrossRef]
13. Kantamneni, A.; Brown, L.E.; Parker, G.; Weaver, W.W. Survey of Multi-Agent Systems for Microgrid Control. *Eng. Appl. Artif. Intell.* **2015**, *45*, 192–203. [CrossRef]

14. Olfati-Saber, R.; Fax, J.A.; Murray, R.M. Consensus and Cooperation in Networked Multi-Agent Systems. *Proc. IEEE* **2007**, *95*, 215–233. [CrossRef]
15. de Witt, C.S.; Foerster, J.N.; Farquhar, G.; Torr, P.H.S.; Boehmer, W.; Whiteson, S. Multi-Agent Common Knowledge Reinforcement Learning. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2019; Volume 32.
16. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*; Proceedings of Machine Learning Research (PMLR): Cambridge, MA, USA, 2017; Volume 54, pp. 1273–1282.
17. Kairouz, P.; McMahan, H.B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A.N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* **2021**, *14*, 1–210. [CrossRef]
18. Liao, W.; Bak-Jensen, B.; Pillai, J.R.; Wang, Y.; Wang, Y. A Review of Graph Neural Networks and Their Applications in Power Systems. *J. Mod. Power Syst. Clean Energy* **2022**, *10*, 345–360. [CrossRef]
19. Bao, J.; Basu, P.; Dean, M.; Partridge, C.; Swami, A.; Leland, W.; Hendler, J.A. Towards a Theory of Semantic Communication. In Proceedings of the IEEE Network Science Workshop (NSW), West Point, NY, USA, 22–24 June 2011; pp. 110–117. [CrossRef]
20. Strinati, E.C.; Barbarossa, S. 6G Networks: Beyond Shannon Towards Semantic and Goal-Oriented Communications. *Comput. Netw.* **2021**, *190*, 107930. [CrossRef]
21. Lan, Q.; Wen, D.; Zhang, Z.; Zeng, Q.; Chen, X.; Popovski, P.; Huang, K. What Is Semantic Communication? A View on Conveying Meaning in the Era of Machine Intelligence. *J. Commun. Inf. Netw.* **2021**, *6*, 336–371. [CrossRef]
22. Gunduz, D.; Qin, Z.; Aguerri, I.E.; Dhillon, H.S.; Yang, Z.; Yener, A.; Wong, K.K.; Chae, C.B. Beyond Transmitting Bits: Context, Semantics and Task-Oriented Communications. *IEEE J. Sel. Areas Commun.* **2023**, *41*, 5–41. [CrossRef]
23. Xie, H.; Qin, Z.; Li, G.Y.; Juang, B.H. Deep Learning Enabled Semantic Communication Systems. *IEEE Trans. Signal Process.* **2021**, *69*, 2663–2675. [CrossRef]
24. Shi, Y.; Zhou, Y.; Wen, D.; Wu, Y.; Jiang, C.; Letaief, K.B. Task-Oriented Communications for 6G: Vision, Principles, and Technologies. *IEEE Wirel. Commun.* **2023**, *30*, 78–85. [CrossRef]
25. Yates, R.D.; Sun, Y.; Brown, D.R.; Kaul, S.K.; Modiano, E.; Ulukus, S. Age of Information: An Introduction and Survey. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 1183–1210. [CrossRef]
26. Tatikonda, S.; Mitter, S. Control Under Communication Constraints. *IEEE Trans. Autom. Control* **2004**, *49*, 1056–1068. [CrossRef]
27. Diao, X.; Song, Y.; Sahoo, S.; Li, Y. Neuromorphic Event-Driven Semantic Communication in Microgrids. *IEEE Trans. Smart Grid* **2024**, *15*, 4300–4314. [CrossRef]
28. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention Is All You Need. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2017.
29. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, arXiv:2108.07258. [CrossRef]
30. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.T.; Rocktäschel, T.; et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2020.
31. Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; Cao, Y. ReAct: Synergizing Reasoning and Acting in Language Models. In Proceedings of the International Conference on Learning Representations, Kigali, Rwanda, 1–5 May 2023.
32. Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; Scialom, T. Toolformer: Language Models Can Teach Themselves to Use Tools. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2023.
33. Zhang, Y.; Saber, A.M.; Youssef, A.; Kundur, D. Grid-Agent: An LLM-Powered Multi-Agent System for Power Grid Control. *arXiv* **2025**, arXiv:2508.05702. [CrossRef]
34. Wen, Y.; Chen, X. X-GridAgent: An LLM-Powered Agentic AI System for Assisting Power Grid Analysis. *arXiv* **2025**, arXiv:2512.20789. [CrossRef]
35. Kiasari, M.; Aly, H. Agentic Artificial Intelligence for Smart Grids: A Comprehensive Review of Autonomous, Safe and Explainable Control Frameworks. *Energies* **2026**, *19*, 617. [CrossRef]
36. Hatziargyriou, N. (Ed.) *Microgrids: Architectures and Control*; Wiley-IEEE Press: Piscataway, NJ, USA, 2014. [CrossRef]
37. Parhizi, S.; Lotfi, H.; Khodaei, A.; Bahramirad, S. State of the Art in Research on Microgrids: A Review. *IEEE Access* **2015**, *3*, 890–925. [CrossRef]
38. Hirsch, A.; Parag, Y.; Guerrero, J. Microgrids: A Review of Technologies, Key Drivers and Outstanding Issues. *Renew. Sustain. Energy Rev.* **2018**, *90*, 402–411. [CrossRef]
39. Mariam, L.; Basu, M.; Conlon, M.F. Microgrid: Architecture, Policy and Future Trends. *Renew. Sustain. Energy Rev.* **2016**, *64*, 477–489. [CrossRef]

40. IRENA. *Innovation Outlook: Renewable Mini-Grids*; Technical report; International Renewable Energy Agency: Abu Dhabi, United Arab Emirates, 2016.
41. World Bank (ESMAP). *Mini Grids for Half a Billion People: Market Outlook and Handbook for Decision Makers*; Technical report; World Bank, Energy Sector Management Assistance Program: Washington, DC, USA, 2019.
42. Justo, J.J.; Mwasilu, F.; Lee, J.; Jung, J.W. AC-Microgrids Versus DC-Microgrids with Distributed Energy Resources: A Review. *Renew. Sustain. Energy Rev.* **2013**, *24*, 387–405. [CrossRef]
43. Planas, E.; Andreu, J.; Garate, J.I.; de Alegria, I.M.; Ibarra, E. AC and DC Technology in Microgrids: A Review. *Renew. Sustain. Energy Rev.* **2015**, *43*, 726–749. [CrossRef]
44. Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; et al. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. In Proceedings of the First Conference on Language Modeling (COLM), Philadelphia, PA, USA, 7–9 October 2024.
45. Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Zhang, C.; Wang, J.; Wang, Z.; Yau, S.K.S.; Lin, Z.; et al. MetaGPT: Meta Programming for a Multi-Agent Collaborative Framework. In Proceedings of the Twelfth International Conference on Learning Representations (ICLR), Vienna, Austria, 7–11 May 2024.
46. Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. The Rise and Potential of Large Language Model Based Agents: A Survey. *Sci. China Inf. Sci.* **2025**, *68*, 121101. [CrossRef]
47. d'Avila Garcez, A.; Lamb, L.C. Neurosymbolic AI: The 3rd Wave. *Artif. Intell. Rev.* **2023**, *56*, 12387–12406. [CrossRef]
48. Zhao, W.X.; Zhou, K.; Li, J.; Tang, T.; Dong, Z.; Hou, Y.; Zhang, B.; Min, Y.; Zhang, J.; Liu, P.; et al. A Survey of Large Language Models. *Front. Comput. Sci.* **2026**, *20*, 2012627. [CrossRef]
49. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal Policy Optimization Algorithms. *arXiv* **2017**, arXiv:1707.06347. [CrossRef]
50. Barbalho, P.I.N.; Moraes, A.L.; Lacerda, V.A.; Barra, P.H.A.; Fernandes, R.A.S.; Coury, D.V. Reinforcement Learning Solutions for Microgrid Control and Management: A Survey. *IEEE Access* **2025**, *13*, 39782–39799. [CrossRef]
51. Zia, M.F.; Inayat, U.; Anjum, M.J.; Khalid, H.M.; Liaquat, S. Reinforcement Learning in Microgrid Energy Management: Review, Methods and Prospects. *IET Smart Grid* **2026**, *9*, e70097. [CrossRef]
52. Puech, A.; Weiss, J.; Brunschwiler, T.; Hamann, H.F. Optimal Power Grid Operations with Foundation Models. *arXiv* **2024**, arXiv:2409.02148. [CrossRef]
53. Omitaomu, O.A.; Niu, H. Artificial Intelligence Techniques in Smart Grid: A Survey. *Smart Cities* **2021**, *4*, 548–568. [CrossRef]
54. Alcántara, A.; Chatzivasileiadis, S. Trustworthiness Layer for Foundation Models in Power Systems: Application to N-k Contingency Screening. *arXiv* **2026**, arXiv:2602.07995. [CrossRef]
55. NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; Technical report; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. [CrossRef]
56. *ISO/IEC 23894:2023*; Information Technology—Artificial Intelligence—Guidance on Risk Management. ISO/IEC: Geneva, Switzerland, 2023.
57. Sagduyu, Y.E.; Erpek, T.; Yener, A.; Ulukus, S. Will 6G Be Semantic Communications? Opportunities and Challenges from Task-Oriented and Secure Communications to Integrated Sensing. *arXiv* **2024**, arXiv:2401.01531. [CrossRef]
58. OpenAI. GPT-4 Technical Report. *arXiv* **2023**, arXiv:2303.08774. [CrossRef]
59. Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv* **2023**, arXiv:2307.09288. [CrossRef]
60. Liu, X.; David, I. Introduction to Digital Twins for the Smart Grid. *arXiv* **2026**, arXiv:2602.14256.. [CrossRef]
61. Tao, F.; Zhang, H.; Liu, A.; Nee, A.Y.C. Digital Twin in Industry: State-of-the-Art. *IEEE Trans. Ind. Inform.* **2019**, *15*, 2405–2415. [CrossRef]
62. Addo, K.; Kabeya, M.; Ojo, E.E. Neuro-Symbolic AI for Explainable Decision-Making in Autonomous Grid Operations. *Preprints* **2025**, 2025080747. [CrossRef]
63. Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; Mane, D. Concrete Problems in AI Safety. *arXiv* **2016**, arXiv:1606.06565. [CrossRef]
64. Garcia, J.; Fernandez, F. A Comprehensive Survey on Safe Reinforcement Learning. *J. Mach. Learn. Res.* **2015**, *16*, 1437–1480.
65. Hendrycks, D.; Carlini, N.; Schulman, J.; Steinhardt, J. Unsolved Problems in ML Safety. *arXiv* **2022**, arXiv:2109.13916. [CrossRef]
66. Zheng, Y.; Chen, Y.; Qian, B.; Shi, X.; Shu, Y.; Chen, J. A Review on Edge Large Language Models: Design, Execution, and Applications. *ACM Comput. Surv.* **2025**, *57*, 1–35. [CrossRef]
67. Park, J.; Samarakoon, S.; Bennis, M.; Debbah, M. Wireless Network Intelligence at the Edge. *Proc. IEEE* **2019**, *107*, 2204–2239. [CrossRef]
68. Howell, S.; Rezgui, Y.; Hippolyte, J.L.; Jayan, B.; Li, H. Towards the Next Generation of Smart Grids: Semantic and Holonic Multi-Agent Management of Distributed Energy Resources. *Renew. Sustain. Energy Rev.* **2017**, *77*, 193–214. [CrossRef]

69. Billanes, J.D.; Jørgensen, B.N.; Ma, Z. A Framework for Resilient Community Microgrids: Review of Operational Strategies and Performance Metrics. *Energies* **2025**, *18*, 405. [CrossRef]
70. Nair, G.N.; Fagnani, F.; Zampieri, S.; Evans, R.J. Feedback Control Under Data Rate Constraints: An Overview. *Proc. IEEE* **2007**, *95*, 108–137. [CrossRef]
71. Kostina, V.; Hassibi, B. Rate-Cost Tradeoffs in Control. *IEEE Trans. Autom. Control* **2019**, *64*, 4525–4540. [CrossRef]
72. Yuksel, S.; Basar, T. *Stochastic Networked Control Systems: Stabilization and Optimization Under Information Constraints*; Birkhauser: Basel, Switzerland, 2013.
73. Necula, G.C. Proof-Carrying Code. In Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Paris, France, 15–17 January 1997; pp. 106–119. [CrossRef]
74. Guerrero, J.M.; Chandorkar, M.; Lee, T.; Loh, P.C. Advanced Control Architectures for Intelligent Microgrids—Part I: Decentralized and Hierarchical Control. *IEEE Trans. Ind. Electron.* **2013**, *60*, 1254–1262. [CrossRef]
75. Guerrero, J.M.; Loh, P.C.; Lee, T.; Chandorkar, M. Advanced Control Architectures for Intelligent Microgrids—Part II: Power Quality, Energy Storage and AC/DC Microgrids. *IEEE Trans. Ind. Electron.* **2013**, *60*, 1263–1270. [CrossRef]
76. Simpson-Porco, J.W.; Shafiee, Q.; Dorfler, F.; Vasquez, J.C.; Guerrero, J.M.; Bullo, F. Secondary Frequency and Voltage Control of Islanded Microgrids via Distributed Averaging. *IEEE Trans. Ind. Electron.* **2015**, *62*, 7025–7038. [CrossRef]
77. Dorfler, F.; Simpson-Porco, J.W.; Bullo, F. Breaking the Hierarchy: Distributed Control and Economic Optimality in Microgrids. *IEEE Trans. Control Netw. Syst.* **2016**, *3*, 241–253. [CrossRef]
78. Zhao, J.; Dorfler, F. Distributed Control and Optimization in DC Microgrids. *Automatica* **2015**, *61*, 18–26. [CrossRef]
79. Bidram, A.; Lewis, F.L.; Davoudi, A. Distributed Control Systems for Small-Scale Power Networks: Using Multiagent Cooperative Control Theory. *IEEE Control Syst. Mag.* **2014**, *34*, 56–77. [CrossRef]
80. Rocabert, J.; Luna, A.; Blaabjerg, F.; Rodriguez, P. Control of Power Converters in AC Microgrids. *IEEE Trans. Power Electron.* **2012**, *27*, 4734–4749. [CrossRef]
81. Pogaku, N.; Prodanovic, M.; Green, T.C. Modeling, Analysis and Testing of Autonomous Operation of an Inverter-Based Microgrid. *IEEE Trans. Power Electron.* **2007**, *22*, 613–625. [CrossRef]
82. Schiffer, J.; Zonetti, D.; Ortega, R.; Stankovic, A.M.; Sezi, T.; Raisch, J. A Survey on Modeling of Microgrids—From Fundamental Physics to Phasors and Voltage Sources. *Automatica* **2016**, *74*, 135–150. [CrossRef]
83. Brabandere, K.D.; Bolsens, B.; den Keybus, J.V.; Woyte, A.; Driesen, J.; Belmans, R. A Voltage and Frequency Droop Control Method for Parallel Inverters. *IEEE Trans. Power Electron.* **2007**, *22*, 1107–1115. [CrossRef]
84. Vasquez, J.C.; Guerrero, J.M.; Luna, A.; Rodriguez, P.; Teodorescu, R. Adaptive Droop Control Applied to Voltage-Source Inverters Operating in Grid-Connected and Islanded Modes. *IEEE Trans. Ind. Electron.* **2009**, *56*, 4088–4096. [CrossRef]
85. Simpson-Porco, J.W.; Dorfler, F.; Bullo, F. Synchronization and Power Sharing for Droop-Controlled Inverters in Islanded Microgrids. *Automatica* **2013**, *49*, 2603–2611. [CrossRef]
86. Han, H.; Hou, X.; Yang, J.; Wu, J.; Su, M.; Guerrero, J.M. Review of Power Sharing Control Strategies for Islanding Operation of AC Microgrids. *IEEE Trans. Smart Grid* **2016**, *7*, 200–215. [CrossRef]
87. Eghtedarpour, N.; Farjah, E. Power Control and Management in a Hybrid AC/DC Microgrid. *IEEE Trans. Smart Grid* **2014**, *5*, 1494–1505. [CrossRef]
88. Loh, P.C.; Li, D.; Chai, Y.K.; Blaabjerg, F. Autonomous Operation of Hybrid Microgrid With AC and DC Subgrids. *IEEE Trans. Power Electron.* **2013**, *28*, 2214–2223. [CrossRef]
89. Mahmoud, M.S.; Hussain, S.A.; Abido, M.A. Modeling and Control of Microgrid: An Overview. *J. Frankl. Inst.* **2014**, *351*, 2822–2859. [CrossRef]
90. Rawlings, J.B.; Mayne, D.Q.; Diehl, M.M. *Model Predictive Control: Theory, Computation and Design*; Nob Hill Publishing: Madison, WI, USA, 2017.
91. Scattolini, R. Architectures for Distributed and Hierarchical Model Predictive Control—A Review. *J. Process Control* **2009**, *19*, 723–731. [CrossRef]
92. Parisio, A.; Rikos, E.; Tzamalis, G.; Glielmo, L. Use of Model Predictive Control for Experimental Microgrid Optimization. *Appl. Energy* **2014**, *115*, 37–46. [CrossRef]
93. Cominesi, S.R.; Farina, M.; Giulioni, L.; Picasso, B.; Scattolini, R. A Two-Layer Stochastic Model Predictive Control Scheme for Microgrids. *IEEE Trans. Control Syst. Technol.* **2018**, *26*, 1–13. [CrossRef]
94. Wooldridge, M. *An Introduction to MultiAgent Systems*, 2nd ed.; Wiley: Chichester, UK, 2009.
95. Shoham, Y.; Leyton-Brown, K. *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*; Cambridge University Press: Cambridge, UK, 2008.
96. Mesbahi, M.; Egerstedt, M. *Graph Theoretic Methods in Multiagent Networks*; Princeton University Press: Princeton, NJ, USA, 2010.
97. Jennings, N.R. On Agent-Based Software Engineering. *Artif. Intell.* **2000**, *117*, 277–296. [CrossRef]
98. Ren, W.; Beard, R.W. Consensus Seeking in Multiagent Systems Under Dynamically Changing Interaction Topologies. *IEEE Trans. Autom. Control* **2005**, *50*, 655–661. [CrossRef]

99. Nedic, A.; Ozdaglar, A. Distributed Subgradient Methods for Multi-Agent Optimization. *IEEE Trans. Autom. Control* **2009**, *54*, 48–61. [CrossRef]
100. Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.* **2011**, *3*, 1–122. [CrossRef]
101. Yang, T.; Yi, X.; Wu, J.; Yuan, Y.; Wu, D.; Meng, Z.; Hong, Y.; Wang, H.; Lin, Z.; Johansson, K.H. A Survey of Distributed Optimization. *Annu. Rev. Control* **2019**, *47*, 278–305. [CrossRef]
102. Logenthiran, T.; Srinivasan, D.; Khambadkone, A.M.; Aung, H.N. Multiagent System for Real-Time Operation of a Microgrid in Real-Time Digital Simulator. *IEEE Trans. Smart Grid* **2012**, *3*, 925–933. [CrossRef]
103. Pipattanasomporn, M.; Feroze, H.; Rahman, S. Multi-Agent Systems in a Distributed Smart Grid: Design and Implementation. In Proceedings of the IEEE/PES Power Systems Conference and Exposition, Seattle, WA, USA, 15–18 March 2009.
104. Morstyn, T.; Hredzak, B.; Agelidis, V.G. Cooperative Multi-Agent Control of Heterogeneous Storage Devices Distributed in a DC Microgrid. *IEEE Trans. Power Syst.* **2016**, *31*, 2974–2986. [CrossRef]
105. Zheng, Y.; Song, Y.; Hill, D.J.; Zhang, Y. Multiagent System Based Microgrid Energy Management via Asynchronous Consensus ADMM. *IEEE Trans. Energy Convers.* **2018**, *33*, 886–888. [CrossRef]
106. Jacobson, D.; Dickerman, L. Distributed Intelligence: A Critical Piece of the Microgrid Puzzle. *Electr. J.* **2019**, *32*, 10–13. [CrossRef]
107. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. 623–656. [CrossRef]
108. Weaver, W. Recent Contributions to the Mathematical Theory of Communication. In *The Mathematical Theory of Communication*; University of Illinois Press: Champaign, IL, USA, 1949.
109. Carnap, R.; Bar-Hillel, Y. *An Outline of a Theory of Semantic Information*; Research Laboratory of Electronics Technical Report 247; Massachusetts Institute of Technology: Cambridge, MA, USA, 1952.
110. Bourtsoulatze, E.; Kurka, D.B.; Gunduz, D. Deep Joint Source-Channel Coding for Wireless Image Transmission. *IEEE Trans. Cogn. Commun. Netw.* **2019**, *5*, 567–579. [CrossRef]
111. Kurka, D.B.; Gunduz, D. DeepJSCC-f: Deep Joint Source-Channel Coding of Images with Feedback. *IEEE J. Sel. Areas Inf. Theory* **2020**, *1*, 178–193. [CrossRef]
112. Jankowski, M.; Gunduz, D.; Mikolajczyk, K. Wireless Image Retrieval at the Edge. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 89–100. [CrossRef]
113. Shao, J.; Mao, Y.; Zhang, J. Task-Oriented Communication for Multi-Device Cooperative Edge Inference. *IEEE Trans. Wirel. Commun.* **2023**, *22*, 73–87. [CrossRef]
114. Popovski, P.; Trillingsgaard, K.F.; Simeone, O.; Durisi, G. 5G Wireless Network Slicing for eMBB, URLLC and mMTC: A Communication-Theoretic View. *IEEE Access* **2018**, *6*, 55765–55779. [CrossRef]
115. Halpern, J.Y.; Moses, Y. Knowledge and Common Knowledge in a Distributed Environment. *J. ACM* **1990**, *37*, 549–587. [CrossRef]
116. Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.R.; Yao, S. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2023.
117. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models Are Few-Shot Learners. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2020.
118. Park, J.S.; O'Brien, J.; Cai, C.J.; Morris, M.R.; Liang, P.; Bernstein, M.S. Generative Agents: Interactive Simulacra of Human Behavior. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, San Francisco, CA, USA, 29 October–1 November 2023; pp. 1–22. [CrossRef]
119. Foruzan, E.; Soh, L.K.; Asgarpoor, S. Reinforcement Learning Approach for Optimal Distributed Energy Management in a Microgrid. *IEEE Trans. Power Syst.* **2018**, *33*, 5749–5758. [CrossRef]
120. Zeng, P.; Li, H.; He, H.; Li, S. Dynamic Energy Management of a Microgrid Using Approximate Dynamic Programming and Deep Recurrent Neural Network Learning. *IEEE Trans. Smart Grid* **2019**, *10*, 4435–4445. [CrossRef]
121. Ji, Y.; Wang, J.; Xu, J.; Fang, X.; Zhang, H. Real-Time Energy Management of a Microgrid Using Deep Reinforcement Learning. *Energies* **2019**, *12*, 2291. [CrossRef]
122. Du, Y.; Li, F. Intelligent Multi-Microgrid Energy Management Based on Deep Neural Network and Model-Free Reinforcement Learning. *IEEE Trans. Smart Grid* **2020**, *11*, 1066–1076. [CrossRef]
123. Ye, Y.; Wang, H.; Chen, P.; Tang, Y.; Strbac, G. Safe Deep Reinforcement Learning for Microgrid Energy Management in Distribution Networks With Leveraged Spatial-Temporal Perception. *IEEE Trans. Smart Grid* **2023**, *14*, 3759–3775. [CrossRef]
124. Li, T.; Sahu, A.K.; Talwalkar, A.; Smith, V. Federated Learning: Challenges, Methods and Future Directions. *IEEE Signal Process. Mag.* **2020**, *37*, 50–60. [CrossRef]
125. Hosseinalipour, S.; Li, S.; Inaolaji, A.; Malandra, F.; Herrera, L.; Mastronarde, N. Synergies Between Federated Foundation Models and Smart Power Grids: A Dual Paradigm. *IEEE Electrif. Mag.* **2026**, *14*, 57–65. [CrossRef]
126. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*, 2nd ed.; MIT Press: Cambridge, MA, USA, 2018.

127. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-Level Control Through Deep Reinforcement Learning. *Nature* **2015**, *518*, 529–533. [CrossRef] [PubMed]
128. Lillicrap, T.P.; Hunt, J.J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; Wierstra, D. Continuous Control with Deep Reinforcement Learning. In Proceedings of the International Conference on Learning Representations, San Juan, Puerto Rico, 2–4 May 2016.
129. Fujimoto, S.; van Hoof, H.; Meger, D. Addressing Function Approximation Error in Actor-Critic Methods. In *Proceedings of the International Conference on Machine Learning*; Proceedings of Machine Learning Research (PMLR): Cambridge, MA, USA, 2018; Volume 80, pp. 1587–1596.
130. Haarnoja, T.; Zhou, A.; Abbeel, P.; Levine, S. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the International Conference on Machine Learning*; Proceedings of Machine Learning Research (PMLR): Cambridge, MA, USA, 2018; Volume 80, pp. 1861–1870.
131. Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; Mordatch, I. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2017.
132. Foerster, J.N.; Farquhar, G.; Afouras, T.; Nardelli, N.; Whiteson, S. Counterfactual Multi-Agent Policy Gradients. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
133. Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W.M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J.Z.; Tuyls, K.; et al. Value-Decomposition Networks for Cooperative Multi-Agent Learning Based on Team Reward. In Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems, Stockholm, Sweden, 10–15 July 2018; pp. 2085–2087.
134. Rashid, T.; Samvelyan, M.; de Witt, C.S.; Farquhar, G.; Foerster, J.; Whiteson, S. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the International Conference on Machine Learning*; Proceedings of Machine Learning Research (PMLR): Cambridge, MA, USA, 2018; Volume 80, pp. 4295–4304.
135. Hernandez-Leal, P.; Kartal, B.; Taylor, M.E. A Survey and Critique of Multiagent Deep Reinforcement Learning. *Auton. Agents Multi-Agent Syst.* **2019**, *33*, 750–797. [CrossRef]
136. Yang, Q.; Liu, Y.; Chen, T.; Tong, Y. Federated Machine Learning: Concept and Applications. *ACM Trans. Intell. Syst. Technol.* **2019**, *10*, 12:1–12:19.
137. Glaessgen, E.H.; Stargel, D.S. The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles. In Proceedings of the AIAA/ASME/AHS Adaptive Structures Conference, Honolulu, HI, USA, 23–26 April 2012.
138. Fuller, A.; Fan, Z.; Day, C.; Barlow, C. Digital Twin: Enabling Technologies, Challenges and Open Research. *IEEE Access* **2020**, *8*, 108952–108971. [CrossRef]
139. Kritzinger, W.; Karner, M.; Traar, G.; Henjes, J.; Sihn, W. Digital Twin in Manufacturing: A Categorical Literature Review and Classification. *IFAC-PapersOnLine* **2018**, *51*, 1016–1022. [CrossRef]
140. Jones, D.; Snider, C.; Nassehi, A.; Yon, J.; Hicks, B. Characterising the Digital Twin: A Systematic Literature Review. *CIRP J. Manuf. Sci. Technol.* **2020**, *29*, 36–52. [CrossRef]
141. Rasheed, A.; San, O.; Kvamsdal, T. Digital Twin: Values, Challenges and Enablers From a Modeling Perspective. *IEEE Access* **2020**, *8*, 21980–22012. [CrossRef]
142. Islam, M.K.; Aslam, U.; Khan, M.Z.I.; Ebrahimi, S.; Ferdowsi, F.; Carbone, M.A. Power System Digital Twins in Action: What We Learnt and Where We Go Next. *Digit. Twins Appl.* **2025**, *2*, e70015. [CrossRef]
143. Pearl, J. *Causality: Models, Reasoning and Inference*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2009.
144. Peters, J.; Janzing, D.; Scholkopf, B. *Elements of Causal Inference: Foundations and Learning Algorithms*; MIT Press: Cambridge, MA, USA, 2017.
145. Scholkopf, B.; Locatello, F.; Bauer, S.; Ke, N.R.; Kalchbrenner, N.; Goyal, A.; Bengio, Y. Toward Causal Representation Learning. *Proc. IEEE* **2021**, *109*, 612–634. [CrossRef]
146. Grieves, M. *Digital Twin: Manufacturing Excellence Through Virtual Factory Replication*; White paper; Florida Institute of Technology: Melbourne, FL, USA, 2014.
147. Javaid, S.; Kurose, Y.; Kato, T.; Matsuyama, T. Cooperative Distributed Control Implementation of the Power Flow Coloring Over a Nano-Grid With Fluctuating Power Loads. *IEEE Trans. Smart Grid* **2017**, *8*, 342–352. [CrossRef]
148. Javaid, S.; Kaneko, M.; Tan, Y. Safe Operation Conditions of Electrical Power System Considering Power Balanceability Among Power Generators, Loads and Storage Devices. *Energies* **2021**, *14*, 4460. [CrossRef]
149. Ioannou, I.I.; Javaid, S.; Christophorou, C.; Vassiliou, V.; Pitsillides, A.; Tan, Y. A Distributed AI Framework for Nano-Grid Power Management and Control. *IEEE Access* **2024**, *12*, 43350–43377. [CrossRef]
150. Alshiekh, M.; Bloem, R.; Ehlers, R.; Konighofer, B.; Niekum, S.; Topcu, U. Safe Reinforcement Learning via Shielding. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
151. *IEC 62351 Series*; Power Systems Management and Associated Information Exchange—Data and Communications Security. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2023.

152. NIST. *NISTIR 7628 Revision 1: Guidelines for Smart Grid Cybersecurity*; Technical report; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2014.
153. *IEEE 2030.7-2017*; Specification of Microgrid Controllers. IEEE: Piscataway, NJ, USA, 2017.
154. *IEEE 2030.8-2018*; Testing of Microgrid Controllers. IEEE: Piscataway, NJ, USA, 2018.
155. *IEEE 1547-2018*; Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces. IEEE: Piscataway, NJ, USA, 2018.
156. *IEC 61850-7-420*; Communication Networks and Systems for Power Utility Automation—Distributed Energy Resources Logical Nodes. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2021.
157. IEC. *IEC 61970 and IEC 61968 Common Information Model Standards*; International Electrotechnical Commission (IEC): Geneva, Switzerland, 2024.
158. NIST. *NIST SP 800-82 Revision 3: Guide to Operational Technology (OT) Security*; Technical Report; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. [CrossRef]
159. *IEEE 2030.9-2019*; Recommended Practice for the Planning and Design of the Microgrid. IEEE: Piscataway, NJ, USA, 2019.
160. *ISO/IEC 42001:2023*; Information Technology—Artificial Intelligence—Management System. ISO/IEC: Geneva, Switzerland, 2023.
161. Kuznetsova, E.; Li, Y.F.; Ruiz, C.; Zio, E.; Ault, G.; Bell, K. Reinforcement Learning for Microgrid Energy Management. *Energy* **2013**, *59*, 133–146. [CrossRef]
162. Goh, H.H.; Huang, Y.; Lim, C.S.; Zhang, D.; Liu, H.; Dai, W.; Kurniawan, T.A.; Rahman, S. An Assessment of Multistage Reward Function Design for Deep Reinforcement Learning-Based Microgrid Energy Management. *IEEE Trans. Smart Grid* **2022**, *13*, 4300–4311. [CrossRef]
163. Lee, S.; Seon, J.; Sun, Y.G.; Kim, S.H.; Kyeong, C.; Kim, D.I.; Kim, J.Y. Novel Architecture of Energy Management Systems Based on Deep Reinforcement Learning in Microgrid. *IEEE Trans. Smart Grid* **2024**, *15*, 1646–1658. [CrossRef]
164. Chow, Y.; Nachum, O.; Duenez-Guzman, E.; Ghavamzadeh, M. A Lyapunov-Based Approach to Safe Reinforcement Learning. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2018.
165. Xie, H.; Qin, Z.; Tao, X.; Letaief, K.B. Task-Oriented Multi-User Semantic Communications. *IEEE J. Sel. Areas Commun.* **2022**, *40*, 2584–2597. [CrossRef]
166. Uysal, E.; Kaya, O.; Ephremides, A.; Gross, J.; Codreanu, M.; Popovski, P.; Assaad, M.; Liva, G.; Munari, A.; Soret, B.; et al. Semantic Communications in Networked Systems: A Data Significance Perspective. *IEEE Netw.* **2022**, *36*, 233–240. [CrossRef]
167. Niazi, M.A.; Kumar, V.; Aslam, U.; Hassan, S.R.; Lee, K.; Shabbir, N. A Semantic Risk-Aware Optimization Framework for Virtual Power Plant Dispatch Using Large Language Models. *Energies* **2026**, *19*, 2820. [CrossRef]
168. Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. A Survey on Large Language Model Based Autonomous Agents. *Front. Comput. Sci.* **2024**, *18*, 186345. [CrossRef]
169. d'Avila Garcez, A.; Broda, K.; Gabbay, D.M. *Neural-Symbolic Learning Systems: Foundations and Applications*; Springer: Berlin/Heidelberg, Germany, 2002.
170. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated Optimization in Heterogeneous Networks. In Proceedings of the Machine Learning and Systems, Austin, TX, USA, 2–4 March 2020.
171. Javaid, S.; Kato, T.; Matsuyama, T. Power Flow Coloring System Over a Nanogrid With Fluctuating Power Sources and Loads. *IEEE Trans. Ind. Inform.* **2017**, *13*, 3174–3184. [CrossRef]
172. Ioannou, I.; Vassiliou, V.; Christophorou, C.; Pitsillides, A. Distributed Artificial Intelligence Solution for D2D Communication in 5G Networks. *IEEE Syst. J.* **2020**, *14*, 4232–4241. [CrossRef]
173. Ioannou, I.; Christophorou, C.; Vassiliou, V.; Pitsillides, A. A Novel Distributed AI Framework with ML for D2D Communication in 5G/6G Networks. *Comput. Netw.* **2022**, *211*, 108987. [CrossRef]
174. Ioannou, I.; Christophorou, C.; Vassiliou, V.; Lestas, M.; Pitsillides, A. Dynamic D2D Communication in 5G/6G Using a Distributed AI Framework. *IEEE Access* **2022**, *10*, 62772–62799. [CrossRef]
175. Ioannou, I.; Nagaradjane, P.; Vassiliou, V.; Pitsillides, A.; Christophorou, C. *Distributed Artificial Intelligence for 5G/6G Communications: Frameworks with Machine Learning*; CRC Press: Boca Raton, FL, USA, 2024.
176. Ioannou, I.I.; Javaid, S.; Vassiliou, V.; Pitsillides, A.; Tan, Y. Applying an Agent-Based Distributed AI Framework to Forecast Power for the Mini-Grid Stability. In Proceedings of the IEEE International Conference on Consumer Electronics—Taiwan (ICCE-Taiwan), Taichung, Taiwan, 9–11 July 2024; pp. 255–256. [CrossRef]
177. Ioannou, I.; Javaid, S.; Tan, Y.; Vassiliou, V. A Hierarchical Predictive-Adaptive Control Framework for State-of-Charge Balancing in Mini-Grids Using Deep Reinforcement Learning. *Electronics* **2026**, *15*, 61. [CrossRef]
178. Ioannou, I.I.; Javaid, S.; Bezha, M.; Tan, Y.; Nagaoka, N.; Vassiliou, V. PhysGTT: A Physics-Guided Self-Supervised Graph Temporal Transformer for Forecasting Electricity Inconsistencies in Mini-Grids. *Energies* **2026**, *19*, 2262. [CrossRef]
179. Javaid, S.; Kaneko, M.; Tan, Y. An Efficient Testing Scheme for Power-Balanceability of Power System Including Controllable and Fluctuating Power Devices. *Designs* **2020**, *4*, 48. [CrossRef]

180. Javaid, S.; Kaneko, M.; Tan, Y. System Condition for Power Balancing Between Fluctuating and Controllable Devices and Optimizing Storage Sizes. *Energies* **2022**, *15*, 1055. [CrossRef]
181. Muhammad, S.; Javaid, S.; Ioannou, I.; Lim, Y.; Tan, Y. Mitigating Energy Losses Under Incremental Load Variations in Distributed Power-Flow Systems While Ensuring User Comfort. *Energies* **2025**, *18*, 5716. [CrossRef]
182. Mo, Y.; Kim, T.H.J.; Brancik, K.; Dickinson, D.; Lee, H.; Perrig, A.; Sinopoli, B. Cyber-Physical Security of a Smart Grid Infrastructure. *Proc. IEEE* **2012**, *100*, 195–209. [CrossRef]
183. Sridhar, S.; Hahn, A.; Govindarasu, M. Cyber-Physical System Security for the Electric Power Grid. *Proc. IEEE* **2012**, *100*, 210–224. [CrossRef]
184. Humayed, A.; Lin, J.; Li, F.; Luo, B. Cyber-Physical Systems Security—A Survey. *IEEE Internet Things J.* **2017**, *4*, 1802–1831. [CrossRef]
185. Cardenas, A.A.; Amin, S.; Sastry, S. Secure Control: Towards Survivable Cyber-Physical Systems. In Proceedings of the International Conference on Distributed Computing Systems Workshops, Beijing, China, 17–20 June 2008.
186. Ten, C.W.; Liu, C.C.; Manimaran, G. Vulnerability Assessment of Cybersecurity for SCADA Systems. *IEEE Trans. Power Syst.* **2008**, *23*, 1836–1846. [CrossRef]
187. Lopes, J.A.P.; Moreira, C.L.; Madureira, A.G. Defining Control Strategies for MicroGrids Islanded Operation. *IEEE Trans. Power Syst.* **2006**, *21*, 916–924. [CrossRef]
188. Piagi, P.; Lasseter, R.H. Autonomous Control of Microgrids. In Proceedings of the IEEE Power Engineering Society General Meeting, Montreal, QC, Canada, 18–22 June 2006.
189. Majumder, R. Some Aspects of Stability in Microgrids. *IEEE Trans. Power Syst.* **2013**, *28*, 3243–3252. [CrossRef]
190. Shafiee, Q.; Guerrero, J.M.; Vasquez, J.C. Distributed Secondary Control for Islanded Microgrids—A Novel Approach. *IEEE Trans. Power Electron.* **2014**, *29*, 1018–1031. [CrossRef]
191. Bidram, A.; Davoudi, A.; Lewis, F.L.; Guerrero, J.M. Distributed Cooperative Secondary Control of Microgrids Using Feedback Linearization. *IEEE Trans. Power Syst.* **2013**, *28*, 3462–3470. [CrossRef]
192. Lu, X.; Yu, X.; Lai, J.; Guerrero, J.M.; Zhou, H. Distributed Secondary Voltage and Frequency Control for Islanded Microgrids With Uncertain Communication Links. *IEEE Trans. Ind. Inform.* **2017**, *13*, 448–460. [CrossRef]
193. Sun, Q.; Li, X.; Zhao, J.; Zhu, J.; Wu, Z.; Gao, T.; Yuan, Y. Joint Planning of Transportable and Flexible Resources in Distribution Systems for Renewable Integration and Resilience Enhancement. *Sustain. Energy Grids Netw.* **2026**, *46*, 102292. [CrossRef]
194. Albadi, M.H.; El-Saadany, E.F. A Summary of Demand Response in Electricity Markets. *Electr. Power Syst. Res.* **2008**, *78*, 1989–1996. [CrossRef]
195. Siano, P. Demand Response and Smart Grids—A Survey. *Renew. Sustain. Energy Rev.* **2014**, *30*, 461–478. [CrossRef]
196. Mohsenian-Rad, A.H.; Wong, V.W.S.; Jatskevich, J.; Schober, R.; Leon-Garcia, A. Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Trans. Smart Grid* **2010**, *1*, 320–331. [CrossRef]
197. Saad, W.; Han, Z.; Poor, H.V.; Basar, T. Game-Theoretic Methods for the Smart Grid: An Overview of Microgrid Systems, Demand-Side Management and Smart Grid Communications. *IEEE Signal Process. Mag.* **2012**, *29*, 86–105. [CrossRef]
198. Vazquez-Canteli, J.R.; Nagy, Z. Reinforcement Learning for Demand Response: A Review of Algorithms and Modeling Techniques. *Appl. Energy* **2019**, *235*, 1072–1089. [CrossRef]
199. Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H.B.; Patel, S.; Ramage, D.; Segal, A.; Seth, K. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, Dallas, TX, USA, 30 October–3 November 2017; pp. 1175–1191. [CrossRef]
200. Dwork, C.; Roth, A. *The Algorithmic Foundations of Differential Privacy*; Now Publishers: Delft, The Netherlands, 2014.
201. Tushar, W.; Yuen, C.; Saha, T.K.; Morstyn, T.; Chapman, A.C.; Alam, M.J.E.; Hanif, S.; Poor, H.V. Peer-to-Peer Energy Systems for Connected Communities: A Review of Recent Advances and Emerging Challenges. *Appl. Energy* **2021**, *282*, 116131. [CrossRef]
202. Sousa, T.; Soares, T.; Pinson, P.; Moret, F.; Baroche, T.; Sorin, E. Peer-to-Peer and Community-Based Markets: A Comprehensive Review. *Renew. Sustain. Energy Rev.* **2019**, *104*, 367–378. [CrossRef]
203. Wang, H.; Huang, J. Incentivizing Energy Trading for Interconnected Microgrids. *IEEE Trans. Smart Grid* **2018**, *9*, 2647–2657. [CrossRef]
204. Gregoratti, D.; Matamoros, J. Distributed Energy Trading: The Multiple-Microgrid Case. *IEEE Trans. Ind. Electron.* **2015**, *62*, 2551–2559. [CrossRef]
205. Morstyn, T.; Teytelboym, A.; McCulloch, M.D. Bilateral Contract Networks for Peer-to-Peer Energy Trading. *IEEE Trans. Smart Grid* **2019**, *10*, 2026–2035. [CrossRef]
206. Mengelkamp, E.; Notheisen, B.; Beer, C.; Dauer, D.; Weinhardt, C. A Blockchain-Based Smart Grid: Towards Sustainable Local Energy Markets. *Comput. Sci. Res. Dev.* **2018**, *33*, 207–214. [CrossRef]
207. Munsing, E.; Mather, J.; Moura, S. Blockchains for Decentralized Optimization of Energy Resources in Microgrid Networks. In Proceedings of the IEEE Conference on Control Technology and Applications, Kohala Coast, HI, USA, 27–30 August 2017.

208. Andoni, M.; Robu, V.; Flynn, D.; Abram, S.; Geach, D.; Jenkins, D.; McCallum, P.; Peacock, A. Blockchain Technology in the Energy Sector: A Systematic Review of Challenges and Opportunities. *Renew. Sustain. Energy Rev.* **2019**, *100*, 143–174. [CrossRef]
209. Fallah, A.; Mokhtari, A.; Ozdaglar, A. Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2020.
210. Briggs, C.; Fan, Z.; Andras, P. Federated Learning with Hierarchical Clustering of Local Updates to Improve Training on Non-IID Data. In Proceedings of the International Joint Conference on Neural Networks, Virtual Event, 19–24 July 2020.
211. Bennis, M.; Debbah, M.; Poor, H.V. Ultra-Reliable and Low-Latency Wireless Communication: Tail, Risk, and Scale. *Proc. IEEE* **2018**, *106*, 1834–1853. [CrossRef]
212. Berkenkamp, F.; Turchetta, M.; Schoellig, A.; Krause, A. Safe Model-Based Reinforcement Learning with Stability Guarantees. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2017.
213. Achiam, J.; Held, D.; Tamar, A.; Abbeel, P. Constrained Policy Optimization. In *Proceedings of the International Conference on Machine Learning*; Proceedings of Machine Learning Research (PMLR): Cambridge, MA, USA, 2017; Volume 70, pp. 22–31.
214. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* **2023**, *55*, 1–38. [CrossRef]
215. Brunke, L.; Greeff, M.; Hall, A.W.; Yuan, Z.; Zhou, S.; Panerati, J.; Schoellig, A.P. Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning. *Annu. Rev. Control Robot. Auton. Syst.* **2022**, *5*, 411–444. [CrossRef]
216. Arrieta, A.B.; Diaz-Rodriguez, N.; Ser, J.D.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges Toward Responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [CrossRef]
217. Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; Pedreschi, D. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* **2018**, *51*, 93:1–93:42. [CrossRef]
218. Lundberg, S.M.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2017.
219. Ribeiro, M.T.; Singh, S.; Guestrin, C. Why Should I Trust You? Explaining the Predictions of Any Classifier. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016.
220. Samek, W.; Montavon, G.; Lapuschkin, S.; Anders, C.J.; Muller, K.R. Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications. *Proc. IEEE* **2021**, *109*, 247–278. [CrossRef]
221. Glavic, M.; Fonteneau, R.; Ernst, D. Reinforcement Learning for Electric Power System Decision and Control: Past Considerations and Perspectives. *IFAC-PapersOnLine* **2017**, *50*, 6918–6927. [CrossRef]
222. Uslar, M.; Specht, M.; Rohjans, S.; Trefke, J.; Gonzalez, J.M. *The Common Information Model CIM: IEC 61968/61970 and 62325—A Practical Introduction to the CIM*; Springer: Berlin/Heidelberg, Germany, 2012.
223. *IEEE 2030.5-2018*; Smart Energy Profile Application Protocol. IEEE: Piscataway, NJ, USA, 2018.
224. OpenADR Alliance. *OpenADR 2.0 Profile Specification B Profile*; OpenADR Alliance: Morgan Hill, CA, USA, 2019.
225. OASIS. *MQTT Version 5.0*; OASIS Open: Woburn, MA, USA, 2019.
226. Gungor, V.C.; Sahin, D.; Kocak, T.; Ergut, S.; Buccella, C.; Cecati, C.; Hancke, G.P. Smart Grid Technologies: Communication Technologies and Standards. *IEEE Trans. Ind. Inform.* **2011**, *7*, 529–539. [CrossRef]
227. Yan, Y.; Qian, Y.; Sharif, H.; Tipper, D. A Survey on Smart Grid Communication Infrastructures: Motivations, Requirements and Challenges. *IEEE Commun. Surv. Tutor.* **2013**, *15*, 5–20. [CrossRef]
228. Schwartz, R.; Dodge, J.; Smith, N.A.; Etzioni, O. Green AI. *Commun. ACM* **2020**, *63*, 54–63. [CrossRef]
229. Shokri, R.; Shmatikov, V. Privacy-Preserving Deep Learning. In Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, 12–16 October 2015.
230. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.C.; Yang, Q.; Niyato, D.; Miao, C. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2031–2063. [CrossRef]
231. Liu, Y.; Ning, P.; Reiter, M.K. False Data Injection Attacks Against State Estimation in Electric Power Grids. *ACM Trans. Inf. Syst. Secur.* **2011**, *14*, 13:1–13:33. [CrossRef]
232. Kosut, O.; Jia, L.; Thomas, R.J.; Tong, L. Malicious Data Attacks on the Smart Grid. *IEEE Trans. Smart Grid* **2011**, *2*, 645–658. [CrossRef]
233. Giani, A.; Bitar, E.; Garcia, M.; McQueen, M.; Khargonekar, P.; Poolla, K. Smart Grid Data Integrity Attacks. *IEEE Trans. Smart Grid* **2013**, *4*, 1244–1253. [CrossRef]
234. European Union. *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*; Publications Office of the European Union: Luxembourg, 2024.
235. Rudin, C. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [CrossRef] [PubMed]

236. Sustainable Energy for All. *State of the Global Mini-Grids Market Report 2020*; Technical report; Sustainable Energy for All: Vienna, Austria, 2020.
237. International Energy Agency. *Africa Energy Outlook 2022*; Technical report; International Energy Agency: Paris, France, 2022.
238. Zimmerman, R.D.; Murillo-Sanchez, C.E.; Thomas, R.J. MATPOWER: Steady-State Operations, Planning and Analysis Tools for Power Systems Research and Education. *IEEE Trans. Power Syst.* **2011**, *26*, 12–19. [CrossRef]
239. Chassin, D.P.; Schneider, K.P.; Gerkensmeyer, C.E. GridLAB-D: An Open-Source Power Systems Modeling and Simulation Environment. In Proceedings of the IEEE/PES Transmission and Distribution Conference and Exposition, Chicago, IL, USA, 21–24 April 2008; pp. 1–5. [CrossRef]
240. Dugan, R.C.; Montenegro, D. *The Open Distribution System Simulator (OpenDSS) Reference Guide*; Technical report; Electric Power Research Institute: Washington, DC, USA, 2020.
241. Brown, T.; Horsch, J.; Schlachtberger, D. PyPSA: Python for Power System Analysis. *J. Open Res. Softw.* **2018**, *6*, 4. [CrossRef]
242. Meinecke, S.; Sarajlic, D.; Drauz, S.R.; Klettke, A.; Lauven, L.P.; Rehtanz, C.; Moser, A.; Braun, M. SimBench: A Benchmark Dataset of Electric Power Systems to Compare Innovative Solutions Based on Power Flow Analysis. *Energies* **2020**, *13*, 3290. [CrossRef]
243. IEEE Power and Energy Society Distribution System Analysis Subcommittee. *Distribution Test Feeders*; IEEE: Piscataway, NJ, USA, 2026. Available online: https://cmte.ieee.org/pes-testfeeders/resources/ (accessed on 4 September 2026).
244. Wilson, E.; Parker, A.; Fontanini, A.; Present, E.; Reyna, J.; Adhikari, R.; Bianchi, C.; CaraDonna, C.; Dahlhausen, M.; Kim, J.; et al. *End-Use Load Profiles for the U.S. Building Stock*; Public dataset; National Renewable Energy Laboratory (NREL), Open Energy Data Initiative (OEDI): Golden, CO, USA, 2021. [CrossRef]
245. Baran, M.E.; Wu, F.F. Network Reconfiguration in Distribution Systems for Loss Reduction and Load Balancing. *IEEE Trans. Power Deliv.* **1989**, *4*, 1401–1407. [CrossRef]
246. Hyndman, R.J.; Koehler, A.B. Another Look at Measures of Forecast Accuracy. *Int. J. Forecast.* **2006**, *22*, 679–688. [CrossRef]
247. Lee, J.D.; See, K.A. Trust in Automation: Designing for Appropriate Reliance. *Hum. Factors* **2004**, *46*, 50–80. [CrossRef] [PubMed]