

Edge AI Energy Management for Solar Battery MicrogridsCecep Sarifudin^{*}, Ari Rinaldi²,¹ Teknik Elektro, Fakultas Teknik, Garut University, Indonesia² Teknik Elektro, Fakultas Teknik, Garut University, Indonesiasarifudin_cecep06@gmail.com, arrinaldi109@gmail.com

Received	05 01 2026	Revised	20 01 2026
Accepted	30 01 2026	Published	02 02 2026

Abstract Solar battery microgrids require energy management decisions that remain safe during forecast error, communication loss, changing electricity prices, and battery degradation. Cloud optimization can support planning, but exclusive dependence on remote computation creates latency, availability, privacy, and cybersecurity concerns. This technical review develops a practical architecture for edge artificial intelligence in microgrid energy management. It connects established hierarchical microgrid control with forecasting, constrained optimization, reinforcement learning, and battery health supervision. The proposed design separates millisecond protection and converter control from minute level scheduling, places a deterministic safety filter between learned policies and actuators, and reserves cloud services for model training, fleet analytics, and long horizon planning. A decision matrix compares rule based control, model predictive control, supervised learning assisted optimization, and reinforcement learning according to data requirements, interpretability, computational load, and commissioning risk. The review also proposes a staged implementation path for small facilities and community systems, with measurable gates for data quality, baseline performance, shadow testing, limited authority trials, and continuous monitoring. Unlike prior surveys that catalogue edge-AI techniques across the broad energy sector, this review is scoped specifically to the safety boundary between learned or optimized policies and deterministic converter control in solar battery microgrids, and it translates that boundary into auditable, staged implementation gates. The central conclusion is that edge AI is most defensible when it augments rather than replaces conventional control. Economic objectives must remain subordinate to electrical constraints, battery health, and graceful degradation. The resulting framework provides researchers and practitioners with a testable basis for designing resilient, auditable, and resource aware microgrid energy management systems.

Keywords *microgrid; edge artificial intelligence; energy management system; battery storage; solar photovoltaic; resilient control*

1 Introduction

Microgrids combine distributed generation, storage, controllable loads, and a defined electrical boundary that can operate with or without the utility grid. Their value is not limited to renewable energy integration. A well designed microgrid can reduce demand peaks, maintain critical services, use local energy more efficiently, and support recovery from upstream outages. These benefits depend on coordinated control across very different time scales. Protection and converter loops act rapidly, secondary control restores voltage and frequency, and energy management schedules resources over minutes or hours [1]-[4]. An optimization method that ignores this hierarchy may produce an economically attractive schedule that is electrically infeasible or unsafe.

Solar photovoltaic generation and facility demand are uncertain, while electrochemical storage carries limits that change with temperature, state of charge, age, and operating history. The energy management system therefore faces a sequential decision problem. It must decide when to charge, discharge, curtail, import, export, or defer flexible loads while preserving reserve and respecting power limits. Model predictive control addresses this problem by repeatedly optimizing over a receding horizon, but its quality depends on models, forecasts, and tractable formulations [5]. Data driven methods can improve forecasts or approximate complex policies, yet they introduce new failure modes such as distribution shift and opaque decisions.

Edge computing changes where these functions can run. Instead of streaming all measurements to a remote platform, an industrial computer near the microgrid can execute forecasts, optimization, anomaly detection, and supervisory logic. Edge placement reduces communication latency and preserves basic operation during loss of wide area connectivity [6], [7]; recent domain-specific surveys confirm that these latency and resilience gains carry through to smart-grid and microgrid deployments specifically, not only to generic edge-computing workloads [19], [20]. It can also reduce exposure of high resolution operational data. These advantages do not eliminate the need for a cloud or control center. Long horizon planning, model training, coordinated updates, and fleet comparisons often benefit from centralized resources. The design question is therefore how to divide authority between device, edge, and cloud layers.

This article offers a technical synthesis and implementation framework rather than a claim of new experimental performance. It focuses on solar battery microgrids serving buildings, campuses, small industries, or communities. The contribution is a layered architecture that keeps safety critical control deterministic, allows artificial intelligence to operate within explicit envelopes, and defines evidence required before increasing automated authority. The discussion covers functional allocation, algorithm selection, data engineering, battery aware objectives, validation, cybersecurity, and deployment economics. Together these elements provide a basis for a later field study with reproducible baselines and clearly bounded claims.

Existing reviews of edge intelligence for power systems tend to be broad in scope, surveying applications across wide-area automation, demand response, electric-vehicle charging, and predictive maintenance, or surveying edge-AI algorithms independently of the electrical safety hierarchy they must operate within [19], [21], [22]. Forecasting-focused studies have likewise shown that edge-deployed models can improve microgrid efficiency [21], but typically stop short of specifying how a learned recommendation is authorized to act on live equipment. The present review addresses that gap directly: it does not propose a new forecasting or control algorithm, but instead specifies where a safety filter must sit between any learned or optimized policy and the actuator, what evidence must accumulate before that policy is granted increasing authority, and how this can be expressed as measurable stage gates. This scoping and the accompanying staged-authority framework, rather than any single algorithm, constitute the article's contribution to the literature.

2 Review Methodology

This review synthesizes peer-reviewed literature on hierarchical microgrid control, model predictive control, edge and fog computing, reinforcement learning for energy management, and lithium-ion battery degradation, retrieved from IEEE Xplore, ScienceDirect, and Google Scholar using combinations of the terms *microgrid*, *edge AI*, *edge intelligence*, *energy management system*, *battery degradation*, *reinforcement learning*, and *model predictive control*. Foundational control-theory sources were included regardless of publication date because the hierarchical control structure they establish remains the technical basis for current practice; application-specific sources on edge computing, learning-based control, and battery health were prioritized from the last ten years to reflect the current state of the technology. Sources were included when they reported a reproducible technical method, a comparative evaluation, or a documented failure mode relevant to microgrid energy management; sources limited to marketing material, unrefereed preprints without methodological detail, or single-vendor product descriptions were excluded. Because included studies differ in task, hardware, and evaluation metric, evidence was synthesized narratively rather than pooled statistically, following the same evidentiary approach used in comparable technical-review articles in this journal. Where sources disagreed on the safety implications of a design choice, the more conservative position was adopted and the disagreement was converted into a validation requirement in Sections 5 and 6, rather than resolved by selective citation.

3 Control and Energy Management Foundations

Hierarchical control provides the structural foundation for intelligent energy management. Primary control stabilizes local converter behavior and shares power, often using droop characteristics. Secondary control corrects residual deviations and coordinates distributed units. Tertiary control

manages exchange with the wider grid and considers economic objectives [2]-[4]. The energy management system usually occupies the tertiary and slower supervisory layers. It should issue setpoints, reserve targets, or operating modes rather than bypassing inverter protection. This separation limits the consequences of a software error and permits each layer to be verified with methods appropriate to its time scale. Broader reviews reinforce the need to match microgrid architecture, demand response, and control strategy to local operating objectives [14]-[16], [18].

A minimum scheduling model includes power balance, component limits, storage dynamics, import and export constraints, and a cost or utility objective. Practical formulations also include reserve, demand charges, curtailment penalties, generator start constraints, comfort bands, and battery degradation proxies. The appropriate level of detail depends on the decision horizon. A day ahead planner can use an aggregate battery model and hourly forecasts, whereas an intra-hour controller needs finer resolution and updated estimates. Excess model detail is not automatically better because unidentifiable parameters can make optimization brittle. The model should be the simplest representation that preserves the decisions and constraints that matter.

Forecasts affect scheduling but should not be treated as certainties. Solar power forecasts may use irradiance, cloud cover, historical output, and short term sky information. Load forecasts may combine calendar variables, production plans, temperature, and recent measurements. Point forecasts are easy to integrate, but prediction intervals or scenarios better expose uncertainty. A controller can translate uncertainty into reserve requirements, chance constraints, or robust bounds. The objective is not to eliminate forecast error. It is to ensure that plausible error does not consume all operational margin or force repeated emergency corrections.

Battery behavior links energy economics to asset life. Frequent cycling can create value, but high temperature, deep cycles, sustained extreme state of charge, and high current may accelerate degradation. Reviews of grid connected storage emphasize that energy management should consider lifetime and application specific duty rather than optimizing energy price alone [12], [13], [17]. A battery aware controller can impose state of charge windows, temperature dependent power limits, throughput costs, or a degradation model. Because degradation estimates are uncertain, the controller should report which assumptions drive dispatch and should support sensitivity analysis before operators accept a policy.

4 Edge AI Architecture

The proposed architecture uses five functional layers. The field layer contains meters, protection relays, inverter controllers, weather sensors, and controllable loads. A deterministic control layer preserves electrical stability and device protection. An edge intelligence layer performs data validation, short horizon forecasting, scheduling, and policy supervision. A site operations layer provides historian, alarms, operator controls, and maintenance workflows. A cloud or enterprise layer supports model development, long horizon studies, backup, and cross site analytics. Each upward dependency must have a defined fallback so that loss of a noncritical layer does not disable lower level safety functions.

Table 1 Functional allocation across the microgrid architecture

Layer	Primary functions	Required fallback
Field devices	Measurement protection converter actuation and breaker status	Local protection remains autonomous
Deterministic control	Primary and secondary control interlocks and operating modes	Validated fixed parameters and safe default mode
Edge intelligence	Data checks forecasts scheduling anomaly detection and safety filtering	Rule based schedule with last known valid configuration
Site operations	Historian alarms operator interface and maintenance records	Local manual control and alarm buffering
Cloud and enterprise	Training fleet analytics long horizon planning and approved updates	No dependency for immediate safe operation

At the edge, a message broker or time series bus can decouple data producers from applications. Measurements require synchronized timestamps, engineering units, quality flags, and source identifiers. The feature pipeline should distinguish raw measurements, cleaned values, derived

features, and model inputs. This lineage is necessary for incident analysis and model reproduction. A local model registry should record version, training interval, input schema, validation results, and approval status. Only signed and approved artifacts should be promoted to active service. The previous stable version should remain available for rollback.

The learned or optimized recommendation must pass through a safety filter before becoming a setpoint. The filter evaluates hard constraints such as charge and discharge power, state of charge bounds, inverter capacity, ramp rate, reserve, breaker status, and islanding mode. It can clip a recommendation, substitute a feasible action, or fall back to a deterministic schedule. This mechanism creates a clear boundary between performance optimization and authority. It also produces an auditable record of attempted actions and rejected constraints, which is more informative than a generic model failure alarm.

Cloud services remain useful but should not sit in the only control path. They can aggregate seasonal data, train forecasting models, evaluate candidate policies, and run planning studies that are too costly for an edge device. Deployment packages can then be transferred to the site during a controlled maintenance window. The edge device verifies integrity, checks schema compatibility, executes a shadow period, and activates the model only after acceptance criteria are met. This division follows the broader edge computing principle of moving latency sensitive inference near data sources while retaining centralized capabilities for resource intensive work [6], [7], [20].

5 Algorithm Options and Selection

Rule based control is the strongest initial baseline because it is understandable, inexpensive, and easy to test. Typical rules preserve a target reserve, limit peak imports, and charge from surplus photovoltaic energy. Their weakness appears when objectives conflict or conditions vary beyond the cases encoded by designers. Model predictive control offers a more systematic alternative by optimizing a future sequence and applying the first action before repeating the calculation [5]. It can express constraints directly and show the predicted consequence of decisions, but it requires solver reliability and reasonably calibrated models.

Supervised learning is often most valuable as an input to optimization rather than as the final controller. Load, solar, tariff, and equipment availability forecasts can improve a conventional scheduler while leaving constraints explicit; edge-deployed forecasting models of this kind have been reported to improve microgrid dispatch efficiency directly [21]. Anomaly detection can prevent questionable sensor values from reaching optimization. This modular arrangement isolates model risk. A forecast can be compared with persistence and seasonal baselines, and a scheduler can be tested with recorded forecasts and with deliberately degraded inputs. Failure of one model need not force replacement of the entire control stack.

Reinforcement learning treats operation as a sequence of state, action, and reward. It can learn policies for uncertain environments without solving an optimization problem at every step, and it can represent long term effects through cumulative reward [8], [9]. Microgrid studies have applied batch and deep reinforcement learning to storage and demand decisions [10], [11]. However, reward design can hide tradeoffs, simulation may not represent rare disturbances, and unconstrained exploration is unacceptable on live electrical equipment. A policy should therefore be trained offline, evaluated against strong baselines, placed behind a safety filter, and initially operated with no actuation authority. Broader reviews of AI for variable renewable-energy systems similarly caution that reported gains from learning-based methods depend heavily on data availability and interpretability requirements that vary by application [22].

Algorithm selection should follow operational requirements rather than novelty. If the site has limited data and a stable tariff, well tuned rules may outperform a complex model after accounting for maintenance burden. If constraints and forecasts dominate, model predictive control is a natural candidate. If nonlinear dynamics or high dimensional decisions make repeated optimization difficult, a learned approximation may help. Reinforcement learning is justified only when sequential effects are material and a credible simulator or historical counterfactual environment exists. The chosen method must be maintainable by the organization that will operate it.

Table 2 Comparison of energy management approaches

Approach	Strength	Main limitation	Recommended role
Rule based	Transparent and easy to commission	Limited coordination under changing conditions	Baseline and fallback
Model predictive control	Explicit objectives constraints and horizon	Model forecast and solver dependence	Primary optimization method
Learning assisted optimization	Improves forecasts and nonlinear estimates	Data drift and pipeline complexity	Modular performance enhancement
Reinforcement learning	Represents sequential decisions and uncertainty	Reward sensitivity and difficult assurance	Offline candidate behind safety filter

6 Data Quality and Validation

A useful data program begins with an asset and signal register. Each point needs a physical meaning, unit, expected range, sample period, timestamp source, calibration responsibility, and retention policy. Measurements should include grid exchange, photovoltaic output, battery power and state estimates, critical feeder demand, breaker state, inverter alarms, ambient conditions, and tariff context. Additional signals are valuable only if they support a defined decision or diagnostic. Collecting large volumes without semantic discipline makes later analysis slower and can conceal silent sensor substitutions or unit errors.

Quality checks should run both at ingestion and before model execution. They include missingness, frozen values, rate of change, cross sensor consistency, clock drift, unit conversion, and operating mode compatibility. A battery cannot simultaneously charge and discharge in a standard aggregate model, and power balance residuals should remain within a tolerance that reflects measurement uncertainty. Instead of deleting questionable values, the pipeline should retain quality flags and document any imputation. Models must be trained and evaluated using the same treatment to avoid discrepancies between laboratory and production behavior.

Validation should use chronological splits because random partitioning can leak seasonal and operational patterns. Forecast models should be compared with naive baselines using multiple horizons and error distributions, not only an average score. Scheduling policies should be evaluated with cost, renewable utilization, peak demand, reserve violations, battery throughput, curtailed energy, and computational time. The primary metric should be specified before comparison. Safety violations cannot be averaged away by economic savings, so feasibility and critical limit breaches should be reported separately.

Testing should progress through simulation, replay, hardware in the loop when available, shadow operation, limited authority, and full supervisory use. Stress scenarios need forecast bias, missing sensors, communication delay, price spikes, unexpected islanding, battery derating, and edge device restart. Each scenario should define an expected safe response. For a learning policy, tests should include input values near and outside the training distribution. Promotion requires not only better central performance but also stable tails, bounded action changes, and successful fallback.

7 Implementation Framework

A staged program reduces technical and organizational risk. The first stage establishes measurement reliability and a deterministic baseline. The second adds forecasting and offline optimization without changing live setpoints. The third deploys shadow recommendations and compares them with operator decisions. The fourth grants limited authority during selected grid connected periods. Later stages can expand operating modes only after evidence accumulates. Each stage should have an owner, acceptance criteria, rollback procedure, and a record of unresolved assumptions.

Table 3 Stage gates for responsible deployment

Stage	Evidence required	Authority
Data readiness	Signal register quality thresholds synchronized time and stable history	No control
Offline benchmark	Chronological evaluation against rule based and optimization baselines	No control
Shadow operation	Live recommendations logged with constraint checks and operator review	No actuation
Limited trial	Stress tests rollback drill and approved operating envelope	Restricted actuation

Stage	Evidence required	Authority
Operational service	Sustained performance monitoring revalidation and change governance	Approved supervisory actuation

The business case should compare whole life cost rather than energy savings alone. Costs include edge hardware, integration, cybersecurity, data engineering, commissioning, model maintenance, operator training, and periodic revalidation. Benefits can include avoided demand charges, reduced fuel use, higher photovoltaic utilization, longer battery life, improved outage readiness, and better diagnostic information. Some benefits are probabilistic. Scenario analysis should therefore show conservative, expected, and favorable cases, and should identify the assumptions that most influence payback.

Governance connects technical performance with operational accountability. A change control process should classify model updates, feature changes, constraint changes, and software patches. Minor retraining with an unchanged schema may follow an automated validation suite, whereas a change to actuation limits requires engineering review. Operators need a concise display of current mode, recommended action, active constraints, confidence indicators, and fallback status. Alarms should distinguish data faults, model faults, infeasible schedules, and equipment faults because they require different responses.

Cybersecurity controls include network segmentation, mutual authentication, least privilege service accounts, signed updates, encrypted communication, protected logs, and backup of configuration and models. The edge device should continue a safe local schedule if external identity or cloud services are unavailable. Remote commands require explicit authorization and should never override relay protection. Physical access, removable media, and vendor support channels are part of the threat surface. Security testing should be included in commissioning and repeated after major architectural changes.

8 Research and Reporting Agenda

Future field studies should report enough context to support comparison. Essential items include component ratings, tariff structure, control interval, forecast horizon, battery limits, degradation assumptions, islanding capability, baseline logic, computing platform, and communications architecture. Public datasets rarely contain all these elements, which makes broad performance claims difficult to interpret. Researchers can improve reproducibility by releasing anonymized time series, configuration files, simulator assumptions, and evaluation code when contractual and privacy conditions permit.

Evaluation should distinguish forecast improvement from control improvement. A better forecast may not reduce cost if the schedule is insensitive within the available reserve, while a robust controller may perform well with modest forecast accuracy. Ablation studies can remove forecast features, degradation costs, safety filters, or uncertainty treatments to show which elements matter. Comparisons should allocate the same information and operational constraints to all methods. Otherwise a learned model may appear superior because the baseline was intentionally simple or denied relevant inputs.

Battery degradation deserves particular attention. Simple throughput prices are transparent but may miss temperature and state dependent aging. High fidelity electrochemical models can be difficult to calibrate and optimize. A useful middle path is to compare several plausible degradation models and report whether the preferred dispatch policy changes. If the decision is stable, model uncertainty is less concerning. If it changes materially, the study should present the tradeoff instead of choosing a single uncertain lifetime estimate.

Another priority is safe adaptation under changing load, generation, and tariffs. Continuous online learning can respond quickly but expands the attack and validation surface. Periodic offline retraining with drift monitoring is easier to govern. Hybrid approaches can adapt forecasts while keeping constraint and dispatch logic fixed. Research should evaluate not only average improvement after adaptation but also detection delay, false drift alarms, rollback reliability, and operator workload. These operational measures determine whether an intelligent controller remains useful after the demonstration period.

9 Limitations

This review is narrative and design-oriented rather than a systematic review with a registered protocol, and it does not claim exhaustive retrieval or formal meta-analysis; heterogeneity of tasks, hardware, and evaluation metrics across the cited studies limits direct quantitative comparison. The proposed architecture and stage-gate framework are conceptual and have not yet been validated on an instrumented field system, so the quantitative claims attributed to specific cited studies (for example, forecasting or communication-overhead improvements) describe the results reported in those original studies and should not be read as performance already demonstrated for the framework proposed here. Because the framework is technology-neutral at several interfaces, implementation choices such as hardware platform, solver, and communications stack are left to the deploying team, and these choices can materially affect cost, latency, and maintainability in ways this review does not quantify.

10 Conclusion

Edge artificial intelligence can strengthen solar battery microgrid energy management when its role is deliberately bounded. The most robust architecture preserves deterministic protection and converter control, performs data validation and short horizon intelligence locally, and uses centralized services for training and planning. Rule based control, model predictive control, supervised learning, and reinforcement learning each have appropriate use cases; none removes the need for explicit electrical constraints and battery health limits. A safety filter, auditable model lifecycle, graceful fallback, and staged authority are core design requirements. The proposed framework turns these principles into measurable implementation gates. Its next step should be validation on a documented microgrid using chronological data, strong baselines, stress scenarios, and reporting that separates economic gains from safety and asset life outcomes.

AUTHOR CONTRIBUTIONS

CRedit: conceptualisation, literature synthesis, framework development, writing, review, and approval of the manuscript.

FUNDING

The authors declare that this research received no funding from any source.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

DATA AVAILABILITY STATEMENT

No original datasets were generated or analyzed in this study, as this article presents a conceptual technical review and design framework.

References

- [1] R. H. Lasseter, MicroGrids, in Proceedings of the 2002 IEEE Power Engineering Society Winter Meeting, vol. 1, pp. 305-308, 2002. <https://doi.org/10.1109/PESW.2002.985003>
- [2] J. M. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuna, and M. Castilla, Hierarchical control of droop-controlled AC and DC microgrids, IEEE Transactions on Industrial Electronics, vol. 58, no. 1, pp. 158-172, 2011. <https://doi.org/10.1109/TIE.2010.2066534>
- [3] A. Bidram and A. Davoudi, Hierarchical structure of microgrids control system, IEEE Transactions on Smart Grid, vol. 3, no. 4, pp. 1963-1976, 2012. <https://doi.org/10.1109/TSG.2012.2197425>
- [4] D. E. Olivares et al., Trends in microgrid control, IEEE Transactions on Smart Grid, vol. 5, no. 4, pp. 1905-1919, 2014. <https://doi.org/10.1109/TSG.2013.2295514>
- [5] A. Parisio, E. Rikos, and L. Glielmo, A model predictive control approach to microgrid operation optimization, IEEE Transactions on Control Systems Technology, vol. 22, no. 5, pp. 1813-1827, 2014. <https://doi.org/10.1109/TCST.2013.2295737>

- [6] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, Edge computing vision and challenges, *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637-646, 2016. <https://doi.org/10.1109/JIOT.2016.2579198>
- [7] Z. Zhou et al., Edge intelligence paving the last mile of artificial intelligence with edge computing, *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1738-1762, 2019. <https://doi.org/10.1109/JPROC.2019.2918951>
- [8] V. Francois-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, An introduction to deep reinforcement learning, *Foundations and Trends in Machine Learning*, vol. 11, nos. 3-4, pp. 219-354, 2018. <https://doi.org/10.1561/22000000071>
- [9] J. R. Vazquez-Canteli and Z. Nagy, Reinforcement learning for demand response a review of algorithms and modeling techniques, *Applied Energy*, vol. 235, pp. 1072-1089, 2019. <https://doi.org/10.1016/j.apenergy.2018.11.002>
- [10] B. Mbuwir, F. Ruelens, F. Spiessens, and G. Deconinck, Battery energy management in a microgrid using batch reinforcement learning, *Energies*, vol. 11, no. 9, article 2466, 2018. <https://doi.org/10.3390/en11092466>
- [11] M. V. Liu, P. M. Reed, D. Gold, G. Quist, and C. L. Anderson, A multiobjective reinforcement learning framework for microgrid energy management, *arXiv:2307.08692*, 2023. <https://doi.org/10.48550/arXiv.2307.08692>
- [12] H. C. Hesse, M. Schimpe, D. Kucevic, and A. Jossen, Lithium-ion battery storage for the grid a review of stationary battery storage system design tailored for applications in modern power grids, *Energies*, vol. 10, no. 12, article 2107, 2017. <https://doi.org/10.3390/en10122107>
- [13] K. Uddin, T. Jackson, W. D. Widanage, G. Chouchelamane, P. A. Jennings, and J. Marco, On the possibility of extending the lifetime of lithium-ion batteries through optimal V2G facilitated by an integrated vehicle and smart-grid system, *Energy*, vol. 133, pp. 710-722, 2017. <https://doi.org/10.1016/j.energy.2017.04.116>
- [14] P. Siano, Demand response and smart grids a survey, *Renewable and Sustainable Energy Reviews*, vol. 30, pp. 461-478, 2014. <https://doi.org/10.1016/j.rser.2013.10.022>
- [15] N. Hatziargyriou, H. Asano, R. Iravani, and C. Marnay, Microgrids, *IEEE Power and Energy Magazine*, vol. 5, no. 4, pp. 78-94, 2007. <https://doi.org/10.1109/MPAE.2007.376583>
- [16] A. Hirsch, Y. Parag, and J. Guerrero, Microgrids a review of technologies key drivers and outstanding issues, *Renewable and Sustainable Energy Reviews*, vol. 90, pp. 402-411, 2018. <https://doi.org/10.1016/j.rser.2018.03.040>
- [17] M. Schimpe et al., Comprehensive modeling of temperature-dependent degradation mechanisms in lithium iron phosphate batteries, *Journal of The Electrochemical Society*, vol. 165, no. 2, pp. A181-A193, 2018. <https://doi.org/10.1149/2.1181714jes>
- [18] F. L. Alvarado-Barrios, C. Alvarez-Arroyo, J. M. Escano, F. M. Gonzalez-Longatt, and J. L. Martinez-Ramos, A review of microgrid architectures and control strategy, *Renewable and Sustainable Energy Reviews*, vol. 144, article 110997, 2021. <https://doi.org/10.1016/j.rser.2021.110997>
- [19] H. B. Gooi, T. Wang, and Y. Tang, Edge intelligence for smart grid: A survey on application potentials, *CSEE Journal of Power and Energy Systems*, vol. 9, no. 5, pp. 1623-1640, 2023.
- [20] P. Zeng, H. Liang, N. Zhang, and C. Song, Recent advances of edge computing for smart grid, *Frontiers in Energy Research*, article 1229000, 2023.
- [21] L. Lv, Z. Wu, L. Zhang, B. B. Gupta, and Z. Tian, An edge-AI based forecasting approach for improving smart microgrid efficiency, *IEEE Transactions on Industrial Informatics*, vol. 18, no. 11, pp. 7946-7954, 2022. <https://doi.org/10.1109/TII.2022.3163137>
- [22] L. A. Yousef, H. Yousef, and L. Rocha-Meneses, Artificial intelligence for management of variable renewable energy systems: A review of current status and future directions, *Energies*, 2023.