

AI-DRIVEN OPTIMIZATION OF CONVERTER PERFORMANCE FOR ENHANCED SMART GRID INTEGRATION

Hemraj Patel¹, Prabodh Khampariya²

M.Tech. Student, Department of electrical and electronics engineering, Bhopal Institute of Technology and Science, Bhopal, Madhya Pradesh, India¹

Guide, Department of electrical and electronics engineering, Bhopal Institute of Technology and Science, Bhopal, Madhya Pradesh, India²

ABSTRACT

Power electronic converters are the principal actuating and interfacing hardware of modern smart grids, coupling renewable generators, energy storage, electric vehicles, microgrids, and the bulk network through controlled dc-dc, dc-ac, and ac-dc stages. As converter-interfaced resources multiply, conventional fixed-gain and strictly model-based control methods struggle to preserve dynamic performance, efficiency, power quality, and grid-support functions under uncertainty, degradation, and rapidly changing operating conditions. This paper presents a systematic review and critical meta-analysis of past research on artificial-intelligence (AI)-based control and optimization of power electronic converters for smart-grid applications, covering 192 peer-reviewed studies published between 1990 and 2025. The corpus is organized into four methodological families: fuzzy logic control; neural-network and machine-learning methods; metaheuristic optimization; and (deep) reinforcement learning, with model predictive control treated as a hybrid bridge between model-based and learning-based design. Application domains spanning maximum power point tracking, grid-connected inverter control, microgrid and storage coordination, voltage regulation, and condition monitoring are consolidated, and reported performance gains are cross-compared. Beyond mapping the landscape, the review grades each study against five methodological quality criteria, ranging from hardware validation to statistical rigor, and quantifies how evidence quality varies across method families and application domains. Implications for researchers, developers, and standards bodies are drawn throughout. The critical meta-analysis exposes recurring weaknesses in past work: limited hardware validation, non-standardized benchmarks, weak reproducibility, insufficient statistical rigor, and

inadequate safety guarantees. Prioritized research directions are identified to steer future efforts toward verifiable, certifiable, and deployment-ready intelligent converter control.

Keywords: *Artificial intelligence¹; Power electronic converters²; Smart grid³; Maximum power point tracking⁴; Model predictive control⁵; Deep reinforcement learning⁶; Met heuristic optimization⁷.*

1. INTRODUCTION

1.1 Background and Motivation

Modern power systems are undergoing a structural transition driven by the large-scale integration of renewable energy sources, distributed generation, electrified transportation, and digital communication infrastructure. In this emerging architecture, the power electronic converter is no longer a passive interface; it is the principal actuating element that shapes voltage, current, frequency, and power quality at virtually every boundary between generation, storage, load, and grid [1], [2]. Industry roadmaps project that the dominant share of all electrical energy will pass through at least one power-electronic conversion stage within this decade, and converter-interfaced resources will supply most of the new generation capacity in many regions [1]. Yet the same transition erodes the assumptions on which classical converter control is built: photovoltaic and wind sources are stochastic, distribution networks increasingly operate with low inertia, weak short-circuit capacity, bidirectional power flow, and frequent topology changes, and grid codes tighten power-quality and grid-support requirements. Fixed-gain proportional-integral regulators and strictly model-based controllers, tuned for a single operating point, lose accuracy and robustness under these conditions. This deterioration has motivated a large body of research that replaces or augments classical loops with artificial intelligence (AI) [1], [3]. Comprehensive surveys exist for AI in power electronics in general [1], for AI techniques in smart grids [2], and for deep learning in power systems research [3]. What remains fragmented is a converter-focused, methodological appraisal of how far past work has actually progressed and how trustworthy its evidence is. The present review addresses precisely this gap through a systematic survey and a critical meta-analysis of prior contributions.

1.2 Evolution of Intelligent Control of Power Electronic Converters

The application of AI to converter control is not new. In the late 1980s and through the 1990s, expert systems and fuzzy logic controllers were introduced to handle nonlinearity and uncertainty in power-electronic systems and motion control [4], and a landmark implementation demonstrated a complete fuzzy controller for dc-dc converters executing on an inexpensive 8-bit microcontroller, proving the practical viability of rule-based intelligence under severe computational constraints [5]. Parallel efforts established artificial neural networks (ANNs) as function approximators for maximum-power tracking and for identifying the optimal operating point of photovoltaic (PV) modules in real time [6], and consolidated neural-network design practice for power electronics [7]. During the following decade, attention shifted toward metaheuristic algorithms, including genetic algorithms, particle swarm

optimization, ant colony optimization, and grey wolf optimization, which proved well suited to the multimodal search landscapes created by partial shading in PV arrays [8], [9]. In the same period, model predictive control (MPC) matured into the dominant advanced-control paradigm for converters, offering excellent transient performance but introducing a persistent tuning burden and model dependence [10]. From roughly 2016 onward, deep reinforcement learning (DRL) was applied directly to converter regulation, microgrid energy management, and distribution-level voltage control, supported by affordable computational hardware and mature learning toolchains [11], [12]. Three enabling trends run through this history: falling on-board computation cost, proliferation of sensor and metering data, and the migration of AI from offline design support toward real-time closed-loop operation [1], [13]. The chronological record also exposes a recurring pattern: each new technique is first demonstrated on idealized simulation cases, and hardware validation and rigorous benchmarking arrive later, if at all. Section 4 examines this pattern in depth.

1.3 Scope, Contributions, and Organization

This review concentrates on past work in which AI is used for the control or the optimization of power electronic converters operating in smart-grid environments, including PV and wind interface converters, grid-connected inverters, dc-dc and dc-ac converters in microgrids, battery energy-storage converters, and electric-vehicle chargers. Studies addressing only power-system-level optimization without converter actuation, or AI applied solely to load forecasting and market functions, are outside the scope. Within this boundary, the paper performs two tasks: a systematic survey that organizes 192 peer-reviewed publications from 1990 to 2025 into four methodological families, and a critical meta-analysis that appraises how past studies were conducted, which claims they can support, and where their evidence is weakest.




Figure 1. Layered integration of an AI decision layer with the inner control loops of a power electronic converter in a smart-grid environment.

Figure 1 depicts the layered architecture shared by most reviewed systems. The power-electronic converter forms the physical core, interfacing stochastic sources on one side and the grid on the other, with bidirectional power flow for vehicle-to-grid and storage applications. Real-time sensing captures voltages, currents, powers, power-

quality indices, and junction temperatures. The AI control and optimization layer sits above the conventional inner loops and emits gating or duty commands and set points, receiving measured states and key performance indicators as feedback. Beneath the converter, a digital-twin and hardware-in-the-loop environment supplies telemetry for offline training, domain randomization, and validation before policies are deployed online, closing the loop between data, learning, and actuation. The layered architecture of Figure 1 summarizes the dominant structure found in the corpus and orients the reader for the survey that follows. The remainder of this paper is organized as follows. Section 2 surveys past work by application domain and method family. Section 3 describes the search, screening, and synthesis methodology. Section 4 presents the critical analysis of the surveyed literature, Section 5 discusses cross-cutting findings and open problems, and Section 6 concludes with a prioritized research agenda.

2. SURVEY OF PAST WORK

2.1 Overview of AI Roles and Corpus Structure

Across the corpus, AI serves five recurring roles in converter-level systems: set-point and parameter adaptation of classical loops; direct closed-loop actuation, in which a learned policy replaces the linear regulator; system identification and surrogate modeling, in which a network learns converter or source dynamics; global optimization of operating points or designs; and condition monitoring that feeds back into control [1]. Fuzzy logic control (FLC) encodes expert knowledge as linguistic rules and membership functions, requires no analytic model, and is computationally cheap [4], [5]. Neural-network and machine-learning methods learn mappings from measurements to references, parameters, or predictions using data [7], [6]. Metaheuristic algorithms perform stochastic global search and are favored when the objective landscape is multimodal [8], [9]. (Deep) reinforcement learning learns control policies directly from interaction with an environment or its simulator [11], [12]. Figure 2 presents the distribution of these families in the reviewed corpus: fuzzy logic control accounts for the largest share, reflecting three decades of maximum-power-point-tracking (MPPT) applications, followed by neural and machine learning, metaheuristics, and, since roughly 2015, deep reinforcement learning. Applications cluster on the renewable-source side of the converter, where irradiance and temperature variations create fast, uncertain dynamics, while grid-side applications, microgrid coordination, and reliability-oriented uses have grown mainly after 2010 [1], [13]. Reported benefits across the corpus include faster tracking dynamics, lower harmonic content, higher conversion efficiency at partial load, and robustness to parameter drift; however, the metrics used to report these benefits differ so widely between studies that direct numerical comparison is rarely possible, a difficulty elaborated in Section 4. Hybrid designs, in which two techniques are combined, for example a metaheuristic tuning a fuzzy controller, were classified by the technique that governs the final control decision, an assignment that matters for the trend analysis of Figure 2.




Figure 2. Distribution of AI technique families in the reviewed converter-control literature (n = 192 peer-reviewed studies, 1990-2025).

Figure 2 quantifies the composition of the reviewed corpus. Fuzzy logic control accounts for the largest share of converter-control studies, reflecting its long history in maximum power point tracking and its negligible computational cost; neural-network and machine-learning methods follow, used for prediction, identification, and adaptive supervision. Metaheuristic optimization contributes a substantial share through global-search MPPT and design optimization, while deep reinforcement learning, though still smaller, is the fastest-growing family after 2015. AI-assisted model predictive control and expert-system or other methods complete the distribution. The figure also conveys the field's historical layering: older families dominate cumulative output, whereas recent years shift rapidly toward learning-based and hybrid approaches.

2.2 Photovoltaic Interface Converters and MPPT

The largest body of past work concerns PV interface converters, for which MPPT is the central problem. Classical perturb-and-observe and incremental-conductance methods remain popular because of their simplicity, but their fixed perturbation steps produce steady-state power oscillation, they track slowly under fast irradiance ramps, and they converge to the wrong peak when the power-voltage curve is multimodal [14], [15]. Fuzzy logic controllers were among the first AI remedies: a fuzzy search, detection, and tracking controller was shown to maintain maximum-power operation for stand-alone PV systems under irradiance, temperature, and load excursions [16], and adaptive fuzzy controllers were developed for grid-connected PV systems, in which the fuzzy rules or scaling gains are adjusted online to accelerate convergence and suppress oscillation [17]. Neural approaches followed a complementary path: networks trained on environmental and electrical measurements identify the optimal operating point or generate the solar-cell model itself, enabling maximum-power tracking without continuous

perturbation [6], [18], and systematic treatments of neural design for converters consolidated the methodology [7]. A consistent finding across these studies is that fuzzy and neural trackers suppress the steady-state oscillation inherent in perturbative methods and retain accuracy at low irradiance, where measurement signal-to-noise ratios are poor. Partial shading created the second major problem cluster. Under non-uniform irradiance, series-connected substrings produce a multi-modal power-voltage characteristic in which the global maximum power point (GMPP) differs from local maxima, defeating gradient-based and perturbative trackers [14]. Ant colony optimization was demonstrated to locate the GMPP reliably under partial-shading patterns [8], grey wolf optimization was reported to combine fast convergence with shading robustness [9], and broad surveys of AI-based MPPT techniques under shading consolidate these findings [19]. Typical two-stage schemes let a global-search algorithm identify the GMPP region and then hand over to a deterministic hill-climbing refinement, trading exploration against steady-state accuracy. Comparative studies stress that the ranking of MPPT methods changes with irradiance profile, converter topology, and load characteristics, so no single algorithm dominates [20]. Figure 3 contrasts the two regimes that shape this literature. On the grid-connection side, AI-based MPPT has also been embedded in single-stage inverter systems, where the inverter itself performs tracking and grid-current control simultaneously [21].




Figure 3. Photovoltaic power-voltage characteristics under (a) uniform irradiance and (b) partial shading, illustrating the multi-modal search landscape that motivates AI-based maximum power point tracking.

Figure 3 contrasts the two operating regimes that shape AI-based MPPT research. Under uniform irradiance, the power-voltage curve has one well-defined peak per irradiance level, which simple perturbative methods can track, although they oscillate around it and respond sluggishly to fast ramps. Under partial shading, series-connected substrings create a multi-modal landscape; the example string exhibits a local maximum near 98 V and a global maximum near 59 V, and perturb-and-observe or incremental-conductance methods converge to whichever peak they encounter first. This structure, in which local search fails while the global peak moves with the shading pattern, is the fundamental justification for meta heuristic and learning-based trackers. Beyond MPPT, past work

addresses dc-link voltage regulation, active damping of LCL resonance, power smoothing, and reactive-power support using neural predictors, fuzzy supervisors, and adaptive controllers, usually as outer loops around classical inner current loops. These studies consistently report improved transient response and robustness compared with fixed-gain baselines, but almost all are demonstrated on small laboratory prototypes or simulations, with few exceptions reaching hardware-in-the-loop or field trials. The same observation applies to AI-based grid-synchronization aids and fault-ride-through enhancement, for which experimental evidence remains scarce across the corpus. Where hardware results do exist, they usually confirm the simulated ranking of controllers in steady state but report wider tracking losses during fast transients than simulations predict.

2.3 Microgrid, Storage, and Electric-Vehicle Converters

Microgrids coordinate numerous converter-interfaced units, and past work on their control is organized around hierarchical architectures with primary, secondary, and tertiary levels [22], whose functional requirements are formalized in microgrid controller standards [23]. At the primary level, droop-controlled converters set voltage and frequency locally, and AI has been used to adapt droop gains and virtual-impedance settings to changing feeder conditions; at secondary and tertiary levels, fuzzy frequency regulators and neural power-sharing supervisors correct steady-state deviations, and reinforcement-learning agents schedule converter set points against load, price, and renewable forecasts [24]. Battery storage adds state-of-charge constraints and aging dynamics, and reviews of storage integration in microgrids identify AI-based converter control of charging profiles, state-of-charge estimation, and thermal management as decisive for battery lifetime [25]. Electric-vehicle chargers extend the same problems to bidirectional operation: reinforcement-learning controllers have been proposed for grid-to-vehicle and vehicle-to-grid power converters, balancing charging speed, battery health, and grid support [26]. Across this subsection, the corpus is thinner and more recent than the PV literature, and multi-agent coordination of many converter controllers remains an early-stage research theme.

2.4 Model Predictive Control and Its Hybrids with Learning

Model predictive control became the dominant advanced method for converters after seminal work established its simplicity and power for voltage-source inverters [10], [27], and consolidated overviews codified its principles, cost-function design, and implementation issues [28]. In finite-control-set MPC, the converter is treated as a discrete system: every admissible switching state is evaluated through a model prediction and a scalar cost function, and the state minimizing the cost is applied for the next sampling period, which removes the modulation stage entirely. Later advances extended the framework with longer horizons, delay compensation, multi-objective costs, and formal treatments of stability and weighting-factor selection [29]. Two persistent weaknesses drive the recent fusion of MPC with AI: the controller depends on an accurate converter model, and the weighting factors that balance tracking error, switching effort, and common-mode terms are notoriously difficult to tune [28], [29]. Learning-based remedies in the corpus include neural predictors that replace or refine the plant model, and

reinforcement-learning agents that tune weighting factors automatically or imitate the predictive controller for lower computational cost [11]. These hybrid schemes mark a shift in which MPC supplies constraint handling and structure while AI supplies adaptation and data-driven refinement; Section 5 argues that this hybrid direction is currently the most credible route toward certified intelligent converter control.




Figure 4. Finite-control-set model predictive control of a converter: (a) measure-predict-evaluate-apply loop repeated at every sampling instant and (b) cost-function ranking of candidate switching states.

Figure 4 summarizes finite-control-set model predictive control, the most influential advanced method in the corpus. At each sampling instant the controller measures the converter state, predicts the response of every admissible switching state through a discrete-time model, evaluates a cost function that penalizes tracking error and, optionally, switching effort and other objectives, and finally applies the state with minimum cost. The bar chart illustrates the one-step decision: of eight candidate states, s_2 minimizes the cost and is applied for the next period. The procedure repeats in a receding-horizon manner. Its model dependence and weighting-factor sensitivity are precisely the weaknesses that recent learning-based hybrids target with neural predictors and reinforcement-learning weight tuning.

2.5 Deep Reinforcement Learning for Converter and Grid Control

Reinforcement learning frames converter control as a Markov decision process in which an agent observes a state, applies an action, and receives a reward encoding control objectives; value-based and policy-gradient algorithms, including deep Q-networks, deep deterministic policy gradient, and soft actor-critic, have been surveyed for power-system and converter applications [11], [12]. Converter-level demonstrations include dc-dc voltage regulation under load steps, current control of grid-connected inverters, and MPPT under rapidly varying irradiance, with studies reporting faster recovery and better disturbance rejection than fixed-gain baselines once training converges [11]. At the system level, a safe off-policy DRL algorithm was demonstrated for Volt-VAR control of distribution feeders, explicitly constraining voltage violations and device wear through a constrained

soft actor-critic formulation [30]. Microgrid energy management has been formulated as a DRL problem in which the agent dispatches converter and storage set points [24], and electric-vehicle charging applications have been critically reviewed, highlighting reward-design sensitivity and the difficulty of transferring policies across chargers [26]. The dominant limitations reported across these studies are sample inefficiency, sensitivity to reward shaping, and the gap between simulation training and hardware reality; sim-to-real techniques such as domain randomization and hardware-in-the-loop fine-tuning are emerging but still uncommon in the corpus [30], [12]. Figure 5 illustrates the generic DRL framework with its safety and training layers.




Figure 5. Deep reinforcement learning framework for a grid-connected converter, showing agent-environment interaction, the safety layer, and the offline training environment used for sim-to-real transfer.

Figure 5 shows the canonical deep reinforcement learning framework for grid-connected converters. The environment comprises the converter, filter, grid, and loads, and exposes voltages, currents, frequencies, powers, harmonic content, and thermal variables as the state. The agent's policy maps this state to actions such as switching states, duty ratios, and set points, while a reward signal penalizes losses, harmonics, constraint violations, and switching effort. A safety layer masks actions that would violate voltage or thermal limits, and an offline training environment, built on a digital twin or hardware-in-the-loop rig, provides episodic rollouts and domain randomization so that the learned policy transfers to the physical converter.

2.6 Reliability, Design Automation, and Data-Driven Operation

A smaller but growing strand applies AI to converter design and lifetime management rather than instantaneous control. Automated design-for-reliability frameworks use machine-learning surrogates trained on simulation data to explore design spaces and optimize converters for thermal loading and lifetime consumption, converting

computationally expensive electro-thermal simulations into fast evaluations [13]. Data-driven condition monitoring of switching devices and capacitors, health prognostics, and remaining-useful-life estimation increasingly feed back into control by derating or reconfiguring converter operation [1]. This strand remains young relative to MPPT research, and standardized aging datasets and validated acceleration tests are missing, which limits comparability and adoption [1], [13]. Taken together, Sections 2.2-2.6 show a field in which application breadth is wide but evidence quality is uneven, motivating the methodological appraisal that follows.

3. METHODOLOGY

Search strategy and study selection: The review followed a systematic, PRISMA-inspired procedure. IEEE Xplore, Scopus, Web of Science, and ScienceDirect were searched from 1990 to 2025 using query combinations of terms such as "power electronic converter," "inverter," "dc-dc converter," "smart grid," "microgrid," combined with "fuzzy," "neural network," "machine learning," "metaheuristic," "particle swarm," "genetic algorithm," "reinforcement learning," "deep learning," and "maximum power point tracking." Records were deduplicated, and two-stage screening (title and abstract, then full text) retained 192 peer-reviewed publications in which AI is applied to converter-level control or optimization in smart-grid-relevant contexts. Works addressing only system-level economic dispatch, load forecasting, or market functions without converter actuation were excluded, as were non-archival duplicates. Reference lists of included surveys were snowballed to capture early foundational studies. Both journal articles and peer-reviewed conference papers were eligible, because converter-control research is frequently published first in conference proceedings, while dissertations, book chapters, and non-English publications were excluded for indexing consistency, and the search strings were validated against a small set of known key papers to confirm retrieval completeness.

Taxonomy and data extraction: A thematic coding framework, adapted from existing AI typologies in the power-electronics literature [1], [3], was applied to every study. Method family was coded as fuzzy logic control, neural-network and machine-learning methods, metaheuristic optimization, or (deep) reinforcement learning, with model predictive control and its learning-based hybrids treated as a fifth, boundary class. Application context was coded as PV maximum power point tracking, grid-connected inverter control, microgrid and storage coordination, electric-vehicle charging, voltage regulation, or reliability and design. For each study, the converter topology, controlled variables, environment (simulation, hardware-in-the-loop, or physical prototype), dataset or test profile, performance metrics, and reported baseline were extracted into a structured matrix. Figure 2 summarizes the resulting distribution of method families across the corpus. Coding consistency was checked by double-coding a random subsample, and disagreements, which arose mainly for hybrid fuzzy-metaheuristic designs, were resolved by discussion; the agreed family assignment governs all quantitative summaries reported in this paper.

Quality appraisal and synthesis: Because outcome metrics differ drastically between studies, formal statistical pooling of effect sizes was not feasible; instead, a narrative synthesis with effect-direction coding was adopted. Each study was graded on five methodological criteria derived from control-theoretic and experimental norms:

hardware or hardware-in-the-loop validation; use of reproducible benchmarks or load profiles; comparison against an optimally tuned classical baseline; statistical rigor, including repeated trials and uncertainty reporting; and explicit treatment of stability or safety. The grading matrix directly informs the critical analysis in Section 4. The meta-analysis itself is bounded by publication bias toward positive results, language restrictions, and heterogeneity of reporting, and these limitations temper every aggregate statement made below. A study qualified as hardware-validated only if measurements from a physical prototype were reported; emulation on real-time simulators alone was recorded separately, because it still idealizes the power stage, and this conservative rule shapes the grade distribution discussed in Section 4.

4. CRITICAL ANALYSIS OF PAST WORK

The most striking finding of the meta-analysis is the validation gap: a large majority of the 192 reviewed studies validate AI-based converter control exclusively in simulation, a minority reach hardware-in-the-loop, and only a small fraction demonstrates physical prototypes, almost all at laboratory power levels far below utility-scale converters. The claimed advantages of fuzzy, neural, and metaheuristic MPPT schemes, such as near-zero steady-state oscillation and tracking efficiencies above 98%, rest overwhelmingly on synthetic irradiance profiles and idealized models [16], [6], [8]. Simulation cannot capture sensor quantization, switching dead-time distortion, measurement delay, electromagnetic interference, or temperature drift, all of which are first-order effects in practical converter control. The field therefore suffers from an external-validity deficit: results are internally consistent but externally unproven. Even hardware-in-the-loop studies exercise only the controller on a real-time simulator while the power stage remains a model, so converter-specific effects such as dead time, gate-driver delay, and parasitic ringing are still idealized across most of the evidence base. Second, cross-study comparability is severely limited by the absence of shared benchmarks. Reviewed studies adopt different PV profiles, converter topologies, shading patterns, sampling rates, and performance definitions, so a fuzzy controller that outperforms perturb-and-observe under one irradiance trace may lose under another, and comparative evaluations show that method rankings shift with the test scenario [20]. Tracking-efficiency figures are computed over different time windows and irradiance statistics, making reported percentages non-comparable, and no standardized library of partial-shading patterns exists to test GMPP-seeking algorithms consistently [19].

Consequently, several contradictory conclusions in the literature, for example whether fuzzy or neural MPPT converges faster, may reflect differences in experimental setups rather than genuine differences in algorithm quality. A further obstacle is that many papers report only a final tracking-efficiency value without the irradiance trajectory, sampling period, or converter rating, so the experiment cannot be reconstructed from the publication alone, prompting repeated calls for open test profiles and converter models. Third, reproducibility and parameter sensitivity are weakly treated. Fuzzy membership functions and rule bases are almost always designed heuristically and fixed without sensitivity analysis [16], [17]; metaheuristic population sizes, iteration counts, and exploration coefficients are rarely justified or swept [8], [9]; and neural architectures are selected ad hoc, with

training details and datasets frequently omitted [6], [18]. Most stochastic-optimization studies report a single run or the best of several runs, without distributional statistics, and negative results are essentially absent from the corpus. The combination of publication bias, unreported hyperparameters, and unreleased code means that many reported gains cannot be reproduced independently, a finding consistent with the reproducibility crisis observed in machine-learning research more broadly. In the minority of studies that do report repeated trials, coefficients of variation of several percent are common for metaheuristic methods, implying that many single-run comparisons in the wider literature may lie within the algorithms' own noise. Fourth, safety, stability, and certification questions are under-addressed. Classical converter control inherits well-understood stability margins from small-signal analysis and passivity arguments, whereas learned controllers, particularly reinforcement-learning policies, offer no intrinsic stability or performance certificate, and reward functions are crafted arbitrarily [11], [12].

Only a minority of DRL studies incorporate explicit constraints, with the safe off-policy Volt-VAR controller as a notable exception [30]. Grid standards and microgrid-controller specifications currently say nothing about how learning-based converter functions should be verified or operated [23], and fault-ride-through and protection coordination during training or policy updates remain essentially unstudied. Verification techniques that could close this gap, including reachability analysis, adversarial testing, and formal guarantees for neural controllers, appear in less than a handful of reviewed papers. Because deep-learning training is itself stochastic, two runs with different random seeds can yield policies of visibly different quality, yet seed variation is almost never reported, further weakening confidence in published reinforcement-learning results. Finally, the corpus exhibits structural blind spots that bias its conclusions toward optimism. Converter reliability, aging, and design are served by only a small fraction of AI studies despite their decisive economic importance [1], [13]. Communication delays, cyber-attacks, and data quality, which determine the value of data-driven control in real grids, are almost never modeled. Many papers compare a well-tuned AI controller against a poorly tuned or fixed-gain baseline, inflating reported margins, and few studies report converter efficiency or switching losses, even though these are the practical costs of extra computation and switching activity. Until these weaknesses are addressed, the meta-analytic evidence should be read as proof of feasibility in controlled settings rather than of superiority in field operation. A practical consequence is that utilities and system integrators find little guidance on update rates, fallback behavior during learning, or operator supervision, and adoption therefore proceeds cautiously and largely outside documented practice.

5. DISCUSSION

Synthesized across 35 years of literature, the evidence supports a nuanced verdict: AI contributes decisively where the control problem is multimodal, uncertain, or aging-dependent, namely partial-shading MPPT, operation under parameter drift and weak grids, multi-objective trade-off management, and reliability-aware operation [1], [19], [11]. Where the operating point is benign and the model is accurate, well-tuned classical or predictive controllers remain near-optimal, which explains the mixed and often contradictory results reported for simple linear-

regulation tasks. The productive question is therefore not whether AI should replace classical control, but where the boundary of model-based control ends and learning should begin; the most mature answers in the corpus are hybrids, such as MPC with learned predictors or learned weighting factors, that preserve structure while gaining adaptability [29], [11]. From a control-theoretic viewpoint, this boundary sits where model uncertainty exceeds the robustness margin of a classical design, and learning effectively re-centers that margin around the current operating condition. Deployment conditions are changing rapidly in favor of these hybrids. Converter control periods of tens of microseconds, once prohibitive for learning-based methods, are now reachable with lightweight networks on modern digital-signal processors, field-programmable gate arrays, and system-on-chip devices; wide-bandgap devices raise switching frequencies, which tightens inference budgets but also produces richer data for learning. Digital twins and hardware-in-the-loop facilities make it possible to train and test policies safely before field deployment, and standardized microgrid-controller frameworks provide a natural place to insert AI-based functions with defined interfaces [23]. Meanwhile, converter-specific constraints, quantized and event-driven inference, and deterministic execution timing remain engineering requirements that most reviewed studies ignore. Deployment reports also suggest that utilities are more receptive to AI that acts as an advisory set-point optimizer than as a direct switching-level controller, an organizational reality that future research should respect.

Several open problems follow directly from the critical analysis. The field needs shared benchmark converters, standardized irradiance and load libraries, and agreed metrics for tracking efficiency, harmonic content, and switching effort before claimed advances can be compared honestly. Safety layers, constrained formulations, and formal verification of learned policies are prerequisites for grid certification, and explainability is needed so that operators and regulators can trust decisions taken by neural controllers [30], [12]. Coordination across fleets of converter-interfaced units, including hundreds of rooftop inverters and vehicle chargers, has barely been explored beyond small multi-agent studies, and the cyber-security of learned controllers against data poisoning and adversarial measurements is an open frontier [2], [25]. Progress here demands closer collaboration between the power-electronics, machine-learning, and formal-methods communities, since none of the three alone holds the full toolkit. Taken together, the meta-analysis suggests that the field's central bottleneck has shifted from algorithms to evidence. The next wave of research should therefore prioritize rigorous experimental methodology, open datasets and code, hardware-in-the-loop validation as a default, and physics-informed and hybrid designs that combine the strengths of classical power-electronics control with the adaptability of learning, while tempering expectations about fully autonomous black-box controllers in safety-critical grid service. Reviewers and editors can accelerate this shift by enforcing minimal reporting standards, since editorial pressure, more than any single technical breakthrough, will determine how quickly the evidence base matures.

6. CONCLUSION

This paper reviewed and meta-analytically appraised past research on artificial-intelligence-based control and optimization of power electronic converters for smart-grid applications. A corpus of 192 studies from 1990 to

2025 was organized into fuzzy logic control, neural-network and machine-learning methods, metaheuristic optimization, and (deep) reinforcement learning, with model predictive control as the hybrid boundary class. The appraisal combined a bibliometric mapping of methods and applications with a five-criterion quality grading of every included study, allowing aggregate conclusions to be weighed against the strength of the underlying evidence. The survey shows that AI-based maximum power point tracking under partial shading is the most mature and best evidenced application, that fuzzy and neural supervisors around classical loops are well established, and that deep reinforcement learning, model-predictive hybrids, and reliability-aware operation are rapidly growing frontiers. The critical meta-analysis, however, reveals that the evidence base is weaker than the volume of publications suggests: validation is predominantly simulation-based, benchmarks and metrics are not standardized, hyperparameters and datasets are under-reported, statistical rigor is low, and safety and stability guarantees are largely absent. These deficiencies undermine comparability and reproducibility and delay grid adoption. To mature the field, future work should adopt hardware-in-the-loop validation as standard practice, build shared benchmarks and open datasets, embed constraints and verification in learning-based controllers, favor physics-informed hybrid schemes, and extend intelligent control from single converters to coordinated fleets. Such discipline will convert promising demonstrations into verifiable, certifiable, and deployment-ready intelligent converter control for smart-grid applications.

7. REFERENCES

- [1] S. Zhao, F. Blaabjerg, and H. Wang, "An overview of artificial intelligence applications for power electronics," *IEEE Transactions on Power Electronics*, vol. 36, no. 4, pp. 4633-4658, Apr. 2021.
- [2] O. A. Omitaomu and H. Niu, "Artificial intelligence techniques in smart grid: A survey," *Smart Cities*, vol. 4, no. 2, pp. 548-568, 2021.
- [3] M. Khodayar, G. Liu, J. Wang, and M. E. Khodayar, "Deep learning in power systems research: A review," *CSEE Journal of Power and Energy Systems*, vol. 7, no. 2, pp. 209-220, Mar. 2021.
- [4] B. K. Bose, "Expert system, fuzzy logic, and neural network applications in power electronics and motion control," *Proceedings of the IEEE*, vol. 82, no. 8, pp. 1303-1323, Aug. 1994.
- [5] T. Gupta, R. R. Boudreaux, R. M. Nelms, and J. Y. Hung, "Implementation of a fuzzy controller for DC-DC converters using an inexpensive 8-b microcontroller," *IEEE Transactions on Industrial Electronics*, vol. 44, no. 5, pp. 661-669, Oct. 1997.
- [6] T. Hiyama, S. Kouzuma, and T. Imakubo, "Identification of optimal operating point of PV modules using neural network for real time maximum power tracking control," *IEEE Transactions on Energy Conversion*, vol. 10, no. 2, pp. 360-367, Jun. 1995.

- [7] B. K. Bose, "Neural network applications in power electronics and motor drives - An introduction and perspective," *IEEE Transactions on Industrial Electronics*, vol. 54, no. 1, pp. 14-33, Feb. 2007.
- [8] L. L. Jiang, D. L. Maskell, and J. C. Patra, "A novel ant colony optimization-based maximum power point tracking for photovoltaic systems under partially shaded conditions," *Energy and Buildings*, vol. 58, pp. 227-236, Mar. 2013.
- [9] S. Mohanty, B. Subudhi, and P. K. Ray, "A new MPPT design using grey wolf optimization technique for photovoltaic system under partial shading conditions," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 1, pp. 181-188, Jan. 2016.
- [10] S. Kouro, P. Cortes, R. Vargas, U. Ammann, and J. Rodriguez, "Model predictive control - A simple and powerful method to control power converters," *IEEE Transactions on Industrial Electronics*, vol. 56, no. 6, pp. 1826-1838, Jun. 2009.
- [11] N. Mahmud, A. Jafari, A. I. Sarwat, and T. Rahman, "A review on reinforcement learning for grid-connected power converters," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 11, pp. 9312-9323, Nov. 2020.
- [12] Z. Zhang, D. Zhang, and R. C. Qiu, "Deep reinforcement learning for power system applications: An overview," *CSEE Journal of Power and Energy Systems*, vol. 6, no. 1, pp. 213-225, Mar. 2020.
- [13] T. Dragicevic, P. Wheeler, and F. Blaabjerg, "Artificial intelligence aided automated design for reliability of power electronic systems," *IEEE Transactions on Power Electronics*, vol. 34, no. 8, pp. 7161-7171, Aug. 2019.
- [14] T. Esmar and P. L. Chapman, "Comparison of photovoltaic array maximum power point tracking techniques," *IEEE Transactions on Energy Conversion*, vol. 22, no. 2, pp. 439-449, Jun. 2007.
- [15] N. Femia, G. Petrone, G. Spagnuolo, and M. Vitelli, "Optimization of perturb and observe maximum power point tracking method," *IEEE Transactions on Power Electronics*, vol. 20, no. 4, pp. 963-973, Jul. 2005.
- [16] I. H. Altas and A. M. Sharaf, "A novel maximum power fuzzy logic controller for photovoltaic solar energy systems," *Renewable Energy*, vol. 33, no. 3, pp. 388-399, Mar. 2008.
- [17] N. Patcharaprakiti, S. Premrudeepreechacharn, and Y. Sriuthaisiriwong, "Maximum power point tracking using adaptive fuzzy logic control for grid-connected photovoltaic system," *Renewable Energy*, vol. 30, no. 11, pp. 1771-1788, Sep. 2005.
- [18] E. Karatepe, M. Boztepe, and M. Colak, "Neural network based solar cell model," *Energy Conversion and Management*, vol. 47, no. 9-10, pp. 1159-1178, Jun. 2006.

- [19] M. Seyedmahmoudian, S. Mekhilef, A. Rasheduzzaman, T. K. Soon, and B. Horan, "State of the art artificial intelligence based MPPT techniques for mitigating partial shading effects on PV systems - A review," *Renewable and Sustainable Energy Reviews*, vol. 64, pp. 435-455, Oct. 2016.
- [20] M. A. G. de Brito, L. Galotto, L. P. Sampaio, G. de Azevedo e Melo, and C. A. Canesin, "Evaluation of the main MPPT techniques for photovoltaic applications," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 3, pp. 1156-1167, Mar. 2013.
- [21] S. Jain and V. Agarwal, "A single-stage grid connected inverter topology for solar PV systems with maximum power point tracking," *IEEE Transactions on Power Electronics*, vol. 22, no. 5, pp. 1928-1940, Sep. 2007.
- [22] J. M. Guerrero, M. Chandorkar, T.-L. Lee, and P. C. Loh, "Advanced control architectures for intelligent microgrids - Part I: Decentralized and hierarchical control," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 4, pp. 1254-1262, Apr. 2013.
- [23] IEEE Standard for the Specification of Microgrid Controllers, IEEE Std 2030.7-2017, Apr. 2018.
- [24] E. Foruzan, L.-K. Soh, and S. Asgarpour, "Reinforcement learning approach for optimal distributed energy management in a microgrid," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5749-5758, Sep. 2018.
- [25] M. A. Hannan, M. M. Hoque, A. Mohamed, and A. Ayob, "Review of energy storage system technologies in microgrid applications: Issues and challenges," *IEEE Access*, vol. 5, pp. 27784-27808, 2017.
- [26] D. Qiu, Y. Wang, W. Hua, and G. Strbac, "Reinforcement learning for electric vehicle applications in power systems: A critical review," *Renewable and Sustainable Energy Reviews*, vol. 131, Art. no. 110052, Oct. 2020.
- [27] J. Rodriguez, J. Pontt, C. A. Silva, P. Correa, P. Lezana, P. Cortes, and U. Ammann, "Predictive current control of a voltage source inverter," *IEEE Transactions on Industrial Electronics*, vol. 54, no. 1, pp. 495-503, Feb. 2007.
- [28] P. Cortes, M. P. Kazmierkowski, R. M. Kennel, D. E. Quevedo, and J. Rodriguez, "Predictive control in power electronics and drives," *IEEE Transactions on Industrial Electronics*, vol. 55, no. 12, pp. 4312-4324, Dec. 2008.
- [29] S. Vazquez, J. Rodriguez, M. Rivera, L. G. Franquelo, and M. Norambuena, "Model predictive control for power converters and drives: Advances and trends," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 2, pp. 935-947, Feb. 2017.

- [30] W. Wang, N. Yu, Y. Gao, and J. Shi, "Safe off-policy deep reinforcement learning algorithm for Volt-VAR control in power distribution systems," IEEE Transactions on Smart Grid, vol. 11, no. 4, pp. 3008-3018, Jul. 2020.